Benchmarking strong-field ionisation with atomic hydrogen

D. Kielpinski[1,2], R. T. Sang[1,2], and I. V. Litvinyuk[2]
[1] *ARC Centre of Excellence for Coherent X-Ray Science and* [2] *Australian Attosecond Science Facility, Centre for Quantum Dynamics, Griffith University, Nathan QLD 4111, Australia*
(Dated: March 27, 2014)

As the simplest atomic system, the hydrogen atom plays a key benchmarking role in laser and quantum physics. Atomic hydrogen is a widely used atomic test system for theoretical calculations of strong-field ionization, since approximate theories can be directly compared to numerical solutions of the time-dependent Schrödinger equation. However, relatively little experimental data is available for comparison to these calculations, since atomic hydrogen sources are difficult to construct and use. We review the existing experimental results on strong-field ionization of atomic hydrogen in multi-cycle and few-cycle laser pulses. Quantitative agreement has been achieved between experiment and theoretical predictions at the 10% uncertainty level, and has been used to develop an intensity calibration method with 1% uncertainty. Such quantitative agreement can be used to certify experimental techniques as being free from systematic errors, guaranteeing the accuracy of data obtained on species other than H. We review the experimental and theoretical techniques that enable these results.

arXiv:1403.6548v1 [physics.atom-ph] 26 Mar 2014

I. INTRODUCTION

Strong-field ionisation is the starting point for experiments that reveal and control the dynamics of atoms, molecules, and solids [1, 2]. The interpretation of such experiments relies on an accurate understanding of the complex, highly nonlinear ionisation dynamics. The experimental data are often quite indirectly related to the physical parameters of interest and can easily show significant discrepancies from theoretical predictions. Retrieving accurate, quantitative information from strong-field experiments therefore requires a careful approach to the assessment and modeling of data.

Atomic hydrogen (H) offers unparalleled opportunities for benchmarking experimental and theoretical techniques in strong-field ionisation. As the simplest bound electronic system, H can be numerically simulated to extremely high precision using the exact time-dependent Schrödinger equation. Careful experiments on H can yield data that quantitatively agrees with the theoretical predictions to within experimental uncertainty. Both the data and the predictions are validated by such agreement. The validated predictions can then be used to calibrate any changes to the experimental parameters, while the validated data can certify the accuracy of the measurement technique for extension to future experiments.

The key benefit offered by experiments on H is *accuracy*, as opposed to precision.[1] Figure 1 illustrates the distinction between the concepts of accuracy and precision. Accuracy refers to the difference between the measured value and the true value of the quantity being measured. Precision refers simply to the difference between the values of repeated measurements. Even if all measurements agree with each other, they may still all differ from the true value. That difference is the systematic error of the measurement, as opposed to the random error associated with lack of precision. The total uncertainty of the measurement is derived from estimates of both systematic and random errors. For H, one can make essentially perfect theoretical predictions of the true values of measurement outcomes. If the measurement and the prediction agree to within the uncertainty of the measurement, we know that the systematic error is insignificant and the measurement is accurate as well as precise.

Figure 1. Distinction between accuracy and precision. The left image shows small random error, i.e., high precision, but large systematic error. The middle image shows low-precision measurements that nevertheless have a small systematic error. The right image shows accurate measurements, which have both low systematic and low random error, and therefore have low uncertainty.

In this review, we discuss the experimental achievements in strong-field ionisation of atomic hydrogen and their relevance to benchmarking of strong-field physics. We summarise the experimental results that have been obtained so far with many-cycle, few-cycle, and carrier-envelope-phase stabilised laser pulses. Soon after the discovery of strong-field ionisation, experiments with many-cycle pulses played a key role in distinguishing resonantly enhanced, perturbative ionisation processes from true nonperturbative strong-field ionisation. In the few-cycle regime, highly accurate comparison of theory and experiment has been achieved with uncertainty at the 10% level. This comparison has been used to measure few-cycle laser peak intensities with an uncertainty of 1%. We detail the experimental and theoretical techniques that enable successful theory-experiment comparisons, as well as the statistical techniques needed for accurate estimation of the experimental uncertainty. While some elements of these techniques are well known to the community, this review brings together all the considerations relevant to state-of-the-art accuracy.

[1] In this review, we use the terminology approved by the Bureau Internationale des Poids et Mesures (BIPM), the intergovernmental body responsible for metrology and precision measurement [3, 4].

II. EXPERIMENTS WITH MANY-CYCLE PULSES

The first experiments on atomic hydrogen in intense laser pulses were performed by Rottke et al. [5] in the group of Welge at Bielefeld. These experiments used 500 fs pulses from an amplified dye laser system, tunable over 596 - 630 nm, to induce multiphoton ionisation. The small pulse bandwidth enabled Rottke et al. to observe electron spectral features with widths below 100 meV. Individual resonantly enhanced multiphoton processes could be identified through the excess kinetic energy imparted by the absorption of the last few photons by the excited atomic state. This interpretation was confirmed by measurements of the photoelectron angular distributions at the resonances. Later, Rottke et al. also observed higher-energy electron peaks (up to 12 eV), which they ascribed to multiphoton ionisation that was resonantly enhanced through the $4f$ excited state [6].

Despite the quite high intensity achieved in these experiments, up to $1.2 \times 10^{14}\,\mathrm{W/cm^2}$, the data obtained by Rottke et al. were readily interpreted within the Floquet theory of multiphoton processes [7]. Since the ponderomotive energy U_p scales as λ^2, the ponderomotive energy in these experiments was about half of that for an equivalent intensity of 800 nm light, so that $U_p \lesssim 4$ eV. The Keldysh parameter for ionisation from the ground state was therefore $\gamma \equiv \sqrt{I_p/(2U_p)} \gtrsim 1.2$, where $I_p = 13.6$ eV is the ionisation energy of H. Naively, one would expect tunnelling physics to be relevant, if not dominant, in this regime. However, for the low electron energies examined in this experiment (< 12 eV), the electron yield was strongly resonantly enhanced by the highly excited $4f$ state and other similarly high-lying states. The binding energy of the $4f$ state is only $E_{4f} = 0.85$ eV, so the final ionisation step was far in the barrier suppression regime (in this case, $I \gtrsim 10^{10}$ W/cm^2). Hence, tunneling physics played a minor role in the interpretation of these results.

Experimental investigations of true strong-field physics in this area were initiated by Paulus et al. [8] in the group of Walther at the Max Planck Institute for Quantum Optics. These experiments examined electron energies up to 40 eV and were the first to observe the ATI plateau in a hydrogen-like system (Figure 2). Paulus et al. actually used deuterium in order to maximise the efficiency of their molecular dissociation source [8]. However, the deuterium atom is still a one-electron, one-nucleus system and the nuclear mass has little effect on strong-field ionisation, so deuterium and hydrogen are expected to exhibit virtually identical photoelectron spectra. Again, visible pulses from an amplified dye laser system were used, with wavelength of 630 nm, but with a much shorter pulse duration of 40 fs. The intensities were similar to those of Rottke et al., up to 0.7×10^{14} W/cm^2.

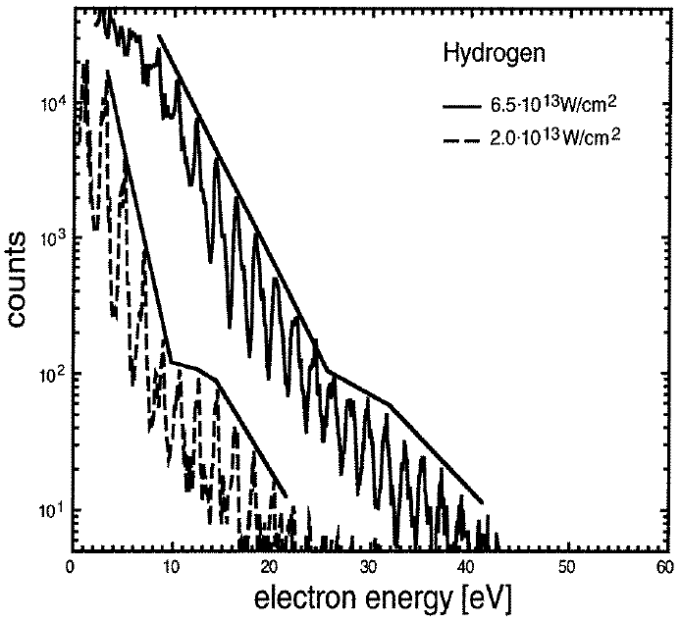

Figure 2. Observation of strong-field ionisation in atomic deuterium [8]. The slanted lines are guides to the eye. A change in slope is clearly observed, indicating the presence of a plateau.

From these experiments, Paulus et al. were able to demonstrate that electrons with energies above ~ 12 eV origi-

nated in strong-field ionisation from the ground state. Plateau features were clearly seen in simulated spectra from numerical solutions of the time-dependent Schrödinger equation, and the energy range of the plateau corresponded to the experimentally observed features. However, as often found in strong-field ionisation experiments, the theoretically predicted strong plateau was significantly obscured by the effect of the inhomogeneous laser intensity, the so-called "focal volume integration". This crucial effect, which is central to any comparison of experiment and theory in strong-field ionisation, is discussed in detail in Section V.

The comparison of Paulus et al's data with the work of Rottke et al. provided further evidence for nonperturbative strong-field dynamics. Resonantly enhanced multiphoton processes were still clearly observed at electron energies of a few eV. However, the plateau peaks were only observed at the higher intensities, indicating the appearance of nonperturbative dynamics as the Keldysh parameter was decreased from $\gamma \approx 2$ to its lowest value of $\gamma = 1.6$. The low-energy electron peaks were also shifted by 0.8 eV relative to the plateau peaks, indicating that they originate from a separate ionisation process.

III. EXPERIMENTS WITH FEW-CYCLE PULSES

Experiments on the ionisation of H with few-cycle pulses commenced in 2009 at the Australian Attosecond Science Facility (AASF) at Griffith University. These experiments have attempted to realise the potential of H to benchmark strong-field ionisation. Few-cycle ionisation and carrier-envelope phase effects have been the main focus, since the physics of these phenomena is highly relevant to attosecond science. The initial goal was to achieve agreement between theory and experiment to within a small, accurately estimated margin of error. Achieving this goal has led to validation of experimental techniques and to novel calibration methods for strong-field experiments.

A. Experimental apparatus

The apparatus used for the AASF experiments is schematically depicted in Figure 3. A few-cycle laser beam intersects an atomic beam of hydrogen. The photoelectrons enter a repeller system, which only passes electrons above a cutoff energy. The high-energy electrons are counted by a channeltron detector. We give a brief review of the experimental design below: further details can be found in [9–11].

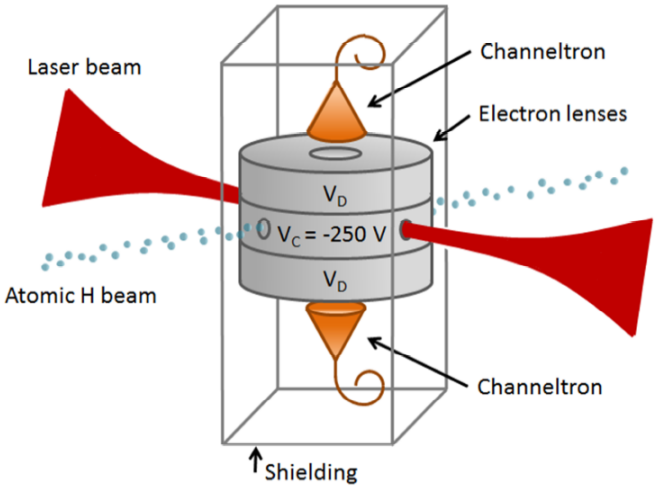

Figure 3. Experimental schematic for few-cycle photoionisation experiments on atomic H.

1. Atomic hydrogen target

The preparation of atomic hydrogen is, in principle, straightforward. H_2 gas is dissociated by electron impact in an electrical discharge or simply by heating to a temperature above the dissociation energy. However, the dissociation process is not particularly efficient, and atomic H readily recombines at most surfaces. For precision measurements, one requires a well-characterised and relatively pure source of H.

The AASF experiments use a radiofrequency (RF) discharge dissociator based on the design of Slevin [12]. The dissociator is essentially identical to those used in the group of Kleppner [13]. H_2 gas at a backing pressure of ~ 2 mbar flows through a Pyrex tube. A discharge is struck within the tube, driven by the high-voltage output of a helical RF resonator operating at 75 MHz. Approximately 10 W of RF power is dissipated in the tube, and water cooling is used to keep the tube temperature below 20 C. The gas that emerges from the discharge tube consists of up to 90% atomic H by number.

The requirements for efficient operation of the H source pose substantial technical difficulties for strong-field ionisation experiments. The outlet of the discharge tube has a 4 mm inner diameter to minimise recombination of the atomic H, much wider than the typical gas-jet nozzle. The gas load of the source is therefore very high, on the order of 5×10^{-3} mbar L/s. Recombination of this high gas load at the chamber walls creates a large background pressure of H_2. At the same time, the backing pressure is approximately $100\times$ lower than in a typical gas-jet experiment, so the atomic density is low.

To isolate the H atoms from the H_2 background, the gas from the source passes from the source chamber, through a few mm aperture into a differential pumping chamber, then through a 0.5 mm aperture into the experiment chamber. The entire system is ~ 2 meters long, much shorter than the mean free path under the vacuum operating conditions, and the apertures are separated by ~ 1 meter. Hence the atoms that emerge into the experiment chamber form a well-collimated atomic beam, while very little of the H_2 gas from the previous chambers escapes into the experiment chamber. However, the number density of the atomic beam is extremely low. Even at the operating pressure of 3×10^{-10} mbar, the background gas in the experiment chamber (mostly H_2O and H_2) makes a significant contribution to the photoionisation signal.

Precision measurements of photoionisation from H require the subtraction of photoionisation signals that originate from the background gas and from undissociated H_2 in the atomic beam. To this end, we perform three sets of measurements: (1) with the discharge turned on, obtaining a signal C_{ON}, (2) with the discharge turned off, obtaining C_{OFF}, (3) with the atomic beam blocked, obtaining C_{BACK}. A pneumatically actuated gate valve between the source chamber and the differential pumping chamber serves to block the atomic beam. The conditions for the three measurements are otherwise made as identical as possible. The signal due to H can then be computed as [9, 10]

$$S_H \propto (C_{\mathrm{ON}} - C_{\mathrm{BACK}}) - (1-\mu)(C_{\mathrm{OFF}} - C_{\mathrm{BACK}}) \tag{1}$$

where $\mu = n_{\mathrm{H}}/(n_{\mathrm{H}} + 2n_{\mathrm{H2}})$ is the mass dissociation fraction of the source. Equation (1) shows that we must know the value of μ in order to make an accurate measurement of S_H.

We determine the value of μ with 10% uncertainty through emission spectroscopy of the RF discharge [14, 15]. In this technique, the light emitted from the discharge is collected and the strengths of the Balmer-α and β lines of atomic hydrogen and several lines in the Fulcher-α band of H_2 are measured with a carefully calibrated monochromator. The ratios of line strengths are compared to a rate-equation model of the radiative and collisional processes in the plasma, allowing the simultaneous determination of μ, the electron temperature of the plasma, and the gas temperature of the plasma. The model has been verified over a wide variety of discharge conditions [15]. Details of our implementation can be found in [10].

2. Laser system

The laser system used in the AASF hydrogen experiments is a Femtopower Compact Pro CE-Phase (Femtolasers GmbH). This system, which is now quite standard for few-cycle experiments, produces pulses of ~ 6 fs duration at 800 nm wavelength. The pulse energy at the target is $\sim 300\,\mu$J, and it is relatively straightforward to focus the pulses to a $1/e^2$ beam radius of $\sim 40\,\mu$m using an off-axis paraboloidal mirror, thus achieving intensities up to 10^{15} W/cm^2.

The laser system consists of a titanium:sapphire (Ti:S) few-cycle oscillator, chirped-pulse Ti:S amplification, and a hollow-fibre compressor. The oscillator produces ~ 6 fs pulses at 80 MHz. Dispersion-compensating mirrors are used to achieve near-zero cavity dispersion in the oscillator. Spectral broadening and difference-frequency generation of the oscillator pulses is used to generate an $f-0f$ carrier-envelope beat signal, which is actively stabilised to 1/4 of the laser repetition rate. After stretching, the oscillator pulses enter a multipass Ti:S amplifier with a repetition rate of 1 kHz and are then recompressed in a prism compressor to ~ 26 fs duration. The amplified pulses undergo spectral broadening in a glass capillary waveguide ("hollow fibre") of ~ 1 m length and 250 μm inner diameter. The spectrally broadened pulses are then recompressed with a dispersion-compensating mirror set to achieve a final pulse duration of $5.5-6.3$ fs, depending on the day-to-day laser stability and the expertise of the operator. A commercial $f-2f$ interferometer is used after the hollow fibre to measure and stabilise the CEP of the amplified laser pulses. The laser beam is focused into the experiment chamber with an off-axis paraboloidal mirror of 750 mm focal length.

At the experimental station, the laser intensity is controlled by transmission through a set of uncoated pellicle membranes on flip mounts. The laser power, and thus intensity, is reduced by Fresnel reflections from each pellicle. The standard technique for controlling the laser power, namely, rotating the laser polarisation with a half-wave plate and then passing the laser through a polariser, appears inadequate for precision measurements with few-cycle pulses. Measurements at the AASF have shown significant deviation from 3D-TDSE predictions when using this technique, whereas the pellicle technique is highly reliable. Optically coated pellicles are found to be suitable for CEP-averaged measurements, but introduce unpredictable changes to the CEP and are therefore unsuitable for comparing CEP-dependent measurements at different intensities. The use of coated glass beamsplitters is also not suitable, since different beamsplitters are found to introduce slightly different wavefront curvature to the beam, thus slightly changing the laser spot size.

The laser beam parameters are carefully measured during the experiment. The laser passes through a low-reflectivity beamsplitter, providing a sample beam for measurement of pulse duration. The beam sample passes through a thickness of fused silica to add a group delay dispersion that is nominally equal to that experienced by the main beam, and its pulse duration is measured by a few-cycle autocorrelator. Another sample beam is picked off after the paraboloidal mirror and sent to a camera-based beam profiler for monitoring of focused spot size and beam quality. The laser power is monitored by a thermal power meter and/or photodiode.

3. Repeller detector

The detection system consists of a set of "repeller" electrodes that pass only the high-energy electrons, two channel electron multipliers ("channeltron") for electron detection, and standard electronics for channeltron signal amplification and discrimination. Figure 4 shows the geometry of the repeller electrodes and the channeltrons. We fix the voltage V_C to -250 V and vary V_T in the range -250 to -350 V, creating a potential barrier that repels the low-energy electrons. The channeltrons are grounded, and the 250 – 350 V acceleration from V_T to the channeltrons ensures that the quantum efficiency of the channeltron detection is nearly independent of electron energy.

The commercial charged-particle optics package `SIMION` was used to simulate electron trajectories for various initial conditions and repeller electrode voltages. These simulations show that the only electrons that reach the channeltron are those with initial energy greater than some cutoff energy E_c. The cutoff energy depends on the "deflection potential" $V_D = V_C - V_T$ as $E_c \approx 0.8(V_D - 5)$, where E_c is measured in eV and V_D is measured in V.

For a given cutoff energy E_c, the number of electrons $Y(E_c)$ detected at one of the channeltrons is approximately proportional to the integral of the electron energy spectrum $S(E)$

$$Y(E_c) \propto \int_{E_c}^{\infty} dE\, S(E) \tag{2}$$

The justification for this approximation and the correction for systematic errors in this detector model are discussed in Section V F.

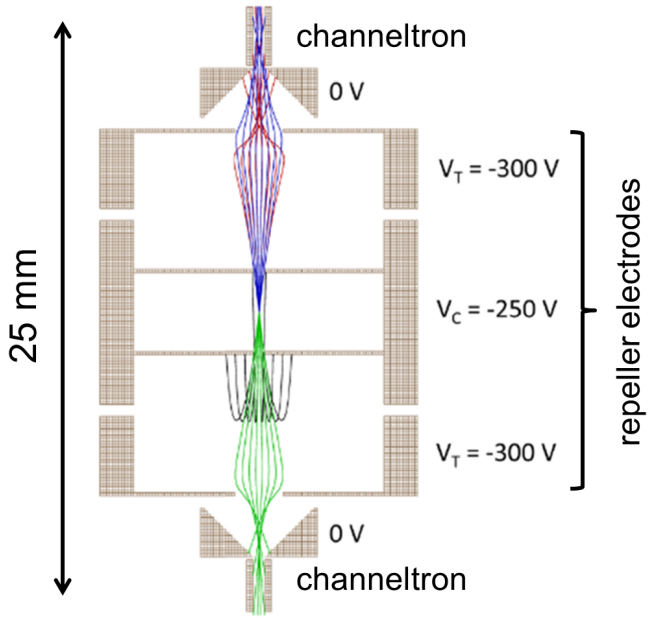

Figure 4. Repeller detector. Electrons originate from photoionisation at the centre of the detector assembly. High-energy electrons overcome the potential barrier between V_C and V_T and are detected at the channeltrons. Electron trajectories (colored lines) are simulated with `SIMION` . Trajectories with initial energy of 25 eV (black lines) do not overcome the barrier, but trajectories with energies of 40 (red), 50 (green), and 80 eV (blue) reach the channeltrons and are detected.

B. Photoelectron measurements

Highly accurate measurements of the few-cycle photoelectron yield $Y(E_c, I)$ as a function of both electron cutoff energy E_c and laser peak intensity I were performed at the AASF in 2010 [9, 10]. For these measurements, the laser pulse duration was 6.3 fs and the laser CEP was made to vary rapidly, thus averaging over a uniform CEP distribution. Over 250 data points were obtained, spanning an order of magnitude in energy (5 - 60 eV) and intensity (1.2 - 5.4 $\times 10^{14}$ W/cm^2). Statistical analysis of the data and estimation of uncertainty was performed using the techniques discussed below, in Section IV. Detailed modelling of the experiment enabled the derivation of realistic predictions of the data from the theoretical simulations, according to the principles of Section V. In fitting the predictions to the experiment, only two fit parameters were used for the entire dataset. Given a theoretical prediction $Y_{\rm pred}(V_D, I)$, we performed a least-squares fit to the entire dataset $Y(V_D, I)$, using the fit function

$$Y_{\rm fit}(V_D, I_{\rm nom}; A, \eta) = A\, Y_{\rm pred}(V_D, \eta I_{\rm nom}) \tag{3}$$

with A and η as fit parameters and $I_{\rm nom}$ as the nominal peak intensity, measured from the laser power, pulse duration, and spot size. The fit parameter η can be interpreted as an overall rescaling of the absolute laser intensity, while A accounts for the absolute quantum efficiency of the detection system and the unknown absolute target density. In the fits, V_D was employed rather than E_c so as to account for any slight deviations from the relation $E_c \approx 0.8(V_T - 5)$. For ease in interpretation, the independent variable was then changed from V_D to E_c using that relation.

Figure 5(a) shows the remarkable agreement between experiment and a prediction based on exact 3D time-dependent Schrödinger equation (3D-TDSE) simulations. The uncertainty in each of the data points was approximately 10%, and again, only two fit parameters were used for all the points, more than 250 in total. Nevertheless, the data and predictions agreed everywhere within the estimated experimental uncertainty. Careful modelling of the experiment was crucial for obtaining this level of agreement. For instance, focal volume averaging effects (Section V E) altered the predicted yield curve by $\sim 50\%$, far more than the experimental uncertainty. For the 3D-TDSE prediction, the intensity rescaling factor η was determined to be 0.94 ± 0.01, well within the usual uncertainty of the method used to measure $I_{\rm nom}$ [16].

In contrast, predictions based on the strong-field approximation (SFA) were sharply ruled out by the data, as shown in Figure 5(b). Since the data was validated by the agreement with 3D-TDSE predictions, the disagreement with

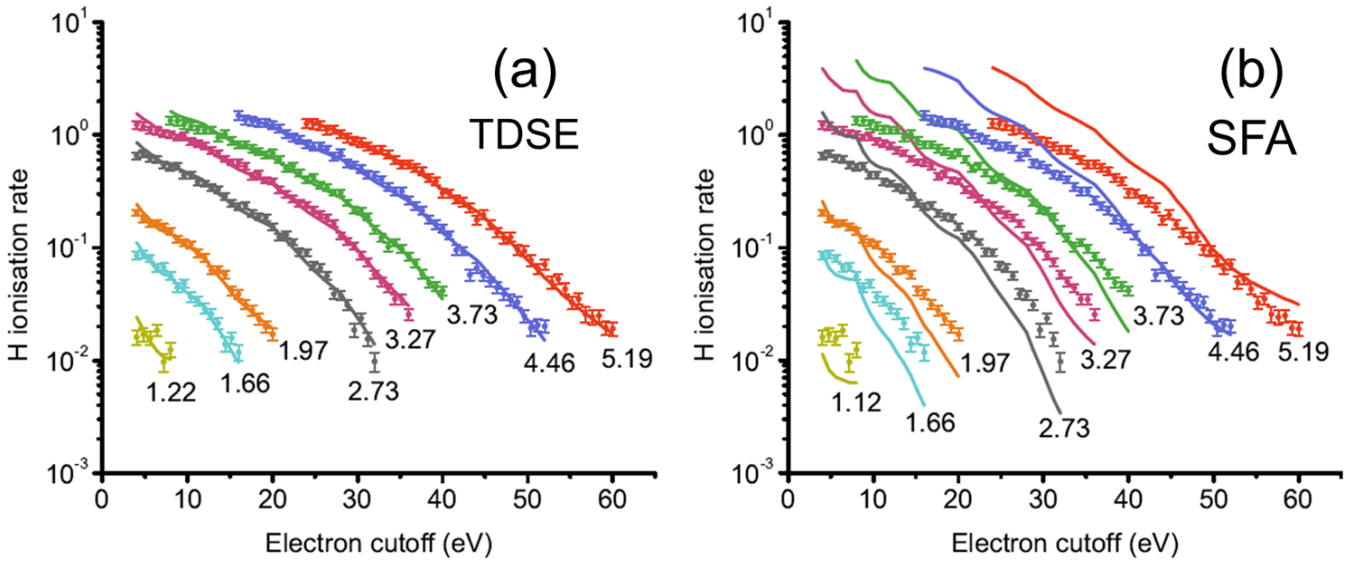

Figure 5. Few-cycle photoelectron yield measurements in H and comparison to theory. Points with error bars: data. Lines: theoretical predictions. The numerical label on each data series gives the intensity in units of 10^{14} W/cm^2. (a) Comparison to TDSE prediction. (b) Comparison to SFA prediction.

the SFA showed that the SFA did not provide adequate predictions in this case. Not only did the prediction fall far outside the experimental uncertainty, but the retrieved value of the intensity factor η was also found to be 1.33 ± 0.05, far outside the uncertainty of $I_{\rm nom}$. As discussed in the Introduction, retrieving reliable physical parameters requires accurate agreement between prediction and experiment. Since the SFA prediction was not adequately accurate, it is not surprising that the SFA retrieved an unreliable value of η.

C. Carrier-envelope phase effects

Experiments to accurately measure CEP effects on the photoelectron yield were undertaken at the AASF in 2011 - 2012 [11]. For these experiments, the laser pulse duration was slightly shorter (5.5 fs). The electron yield $Y_{T,B}(\phi_{\rm CEP}, E_c, I)$ was measured for both top (Y_T) and bottom (Y_B) detectors as a function of CEP, $\phi_{\rm CEP}$, for a range of electron cutoff energies E_c and laser peak intensities I. At each fixed value of E_c and I, the value of $\phi_{\rm CEP}$ was stepped and the yield was measured for each $\phi_{\rm CEP}$ value. The CEP was stepped by changing the insertion of a fused silica wedge pair in the laser beam. Stepping the CEP electronically, via the laser stabilisation system, was found to produce unreliable results.

Following standard practice [17], the CEP effect at a given value of E_c and I was parametrised through the asymmetry, defined by

$$A(\phi_{\rm CEP}) = \frac{Y_T(\phi_{\rm CEP}) - Y_B(\phi_{\rm CEP})}{Y_T(\phi_{\rm CEP}) + Y_B(\phi_{\rm CEP})} \tag{4}$$

$$\approx A_0 \sin(\phi_{\rm CEP} - \phi_0) \tag{5}$$

The asymmetry was found to follow a sinusoidal function of period 2π within the uncertainty of the individual data points, as expressed in the second equality. Fitting the measured $A(\phi_{\rm CEP})$ to this sinusoid yielded the asymmetry amplitude A_0 and phase offset ϕ_0. A map of the amplitude $A_0(E_c, I)$ was constructed, covering a wide range of electron cutoff energies E_c and laser peak intensities I. A map of the phase offset was also constructed [11], but its precision was not high enough to admit useful interpretation and we omit it here.

Figure 6 shows the map of asymmetry amplitude $A_0(E_c, I)$ and compares it to the predictions derived from 3D-TDSE simulations. The data and predictions agreed *with no free parameters* over the entire range of E_c and I. Again,

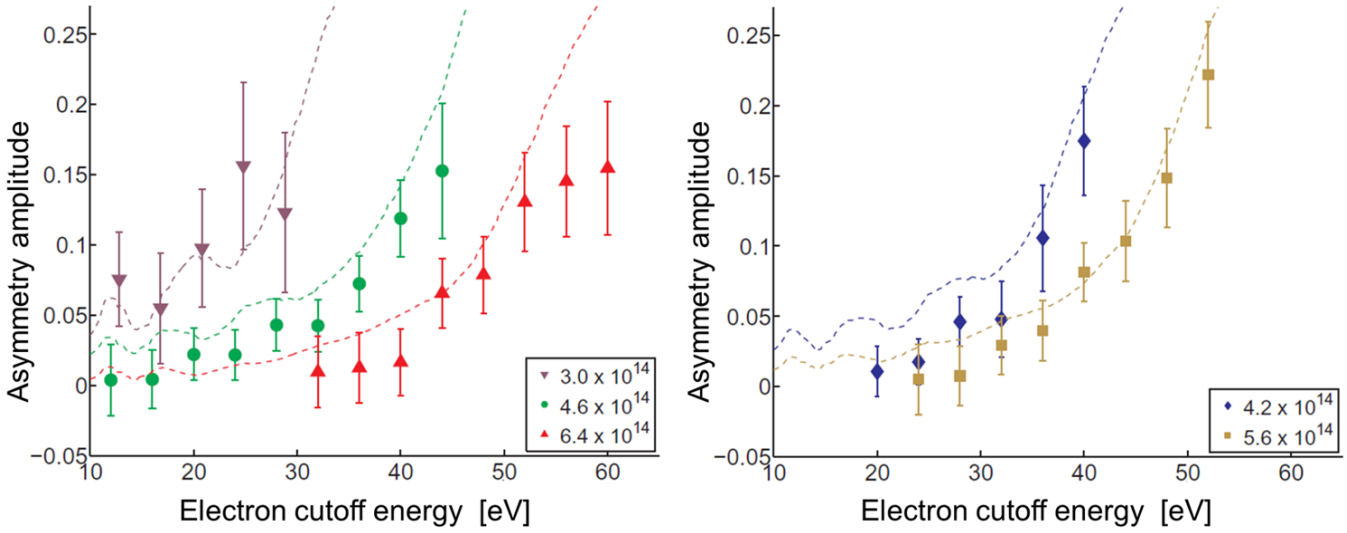

Figure 6. CEP effects on few-cycle photoelectron yield in H. The asymmetry amplitude A_0 is shown as a function of electron energy E_c over a range of intensities (labeled in legends, in units of 10^{14}W/cm^2.

careful modelling of the experiment was essential to accurately predict the observations. In particular, it was found that the 10 meter air path between the laser system and the experiment chamber introduced large fluctuations of the CEP. Measurement of the CEP with an $f-2f$ interferometer at the experiment chamber showed a root-mean-square uncertainty of 850 mrad in the CEP, with a corresponding reduction in the asymmetry amplitude per Section V D. This CEP instability also showed why the simple sinusoidal fit, Eq. (5), was adequate to model the data.

The asymmetry amplitude A_0 was seen to increase rapidly as the electron cutoff energy E_c approached $2U_p$, a trend that could be systematically observed for all intensities. This trend accords with theoretical expectations for the behaviour of high-energy direct electrons [18]. It was not possible to observe the plateau electrons, since the H beam density was too low to provide adequate signal.

D. Intensity calibration standard

The excellent agreement between experiment and theory for H enables highly accurate retrieval of experimental parameters. By obtaining experiment-theory agreement over a wide parameter range, one establishes that the theoretical predictions reliably reflect the experimental conditions. Alternatively, such agreement can be interpreted as showing that the experimental apparatus is accurately calibrated. Later on, if more data are obtained under imperfectly known experimental conditions, these data can be fitted to theory. So long as the fit between theory and data remains within the known uncertainty of the data, the fit coefficients will reliably estimate the experimental parameters.

The laser peak intensity represents a particularly important target for calibration and estimation. Intensity changes of a few percent can radically alter strong-field ionisation [19, 20]. However, the intensity of strong-field laser pulses is difficult to measure directly. One can combine measurements of laser power (via, e.g., power meter), pulse duration (autocorrelator), and spot size (camera-based beam profiler) to obtain an intensity measurement. The uncertainty of this technique is still approximately 10% even when care is taken with the measurements, mostly owing to uncertainty of the laser spot size [16]. Moreover, there may be systematic differences between the measured beam properties and the actual beam properties at the interaction region.

Several methods for *in situ* intensity calibration have been developed, but while these methods offer high precision, they are all vulnerable to uncontrolled systematic errors. The total photoionisation yield may be measured at several estimated intensities and compared to theoretical predictions for the yield [21]. The photoelectron and photoion momentum distribution may be compared to semiclassical theory [22–24]. Finally, the photoelectron spectrum may

be compared to quantitative rescattering theory [25, 26]. In each of these experiments, the data were compared to an approximate theory, but it was not possible to estimate the effect of the approximations on the intensity retrieval, since exact theories are not available for species other than H.

The AASF experiments on H provided intensity calibration with total uncertainty below 1% in the range from 1 to 5 $\times 10^{14}$W/cm^2 [16]. There were no systematic errors arising from theoretical approximations, since the accuracy of the calibration was guaranteed by the agreement between data and 3D TDSE predictions. These experiments also showed that introducing even reasonable and common approximations to the theory can induce errors on the 10% level, comparable to the uncertainty of intensity estimates derived from measuring laser beam properties.

To calibrate the intensity, the AASF team acquired data on the photoelectron yield $Y_{\rm cal}(V_D)$ as a function of deflector voltage V_D, keeping the laser intensity constant. The data was fitted to the function

$$Y_{\rm fit}(V_D; A, I_{\rm TDSE}) = A\, Y_{\rm TDSE}(V_D, I_{\rm TDSE}) \tag{6}$$

where $Y_{\rm TDSE}(V_D, I)$ is the predicted yield based on TDSE simulations. The fit parameters were now $I_{\rm TDSE}$, the laser peak intensity as retrieved by the TDSE, and A, the detection efficiency factor. The fit returned the peak intensity for the dataset.

The intensity calibration procedure should be carefully distinguished from the test of experiment-theory agreement in Section V. In the agreement test, the entire dataset $Y(V_D, I)$, ranging over all intensities, was fitted and found to agree with theory. The global intensity rescaling factor accounted for our uncertainty in measuring the intensity, thus enabling a good match between experiment and theory. For intensity calibration, we *relied* on the global agreement between experiment and theory that was already established in Section V. Given this global agreement, we could be confident that new data $Y_{\rm cal}(V_D)$ would also agree with the TDSE predictions under an appropriate fitting procedure. Agreement of the new data with the TDSE predictions would occur precisely when the fitted intensity correctly estimated the true (unknown) intensity.

Unlike the 3D TDSE, approximate theories were found to be incapable of accurately calibrating the intensity. The fractional systematic error in theory X (e.g., 1D TDSE theory) is given by

$$\epsilon_X = \frac{I_{\rm est,X} - I_{\rm TDSE}}{I_{\rm TDSE}} \tag{7}$$

where $I_{\rm est,X}$ is the intensity estimated using the prediction of theory X. Naturally, ϵ_X varied depending on the actual intensity $I_{\rm TDSE}$. Figure 7 shows the systematic error $\epsilon_X(I_{\rm TDSE})$ for a selection of approximate theories, measured relative to the intensity estimated by the 3D TDSE prediction at 6.3 fs pulse duration. Details of the theoretical methods are given in [16]. The agreement of the two TDSE simulations, for 5.5 and 6.3 fs pulse durations, showed that the intensity calibration technique was relatively insensitive to the pulse duration. However, for all theories except for the 3D TDSE, ϵ_X fell outside the 10% uncertainty range obtained by independent measurements of the laser beam parameters. Only the *in situ* calibration that uses the TDSE could be regarded as an improvement over an intensity estimate based on measurement of the laser beam parameters.

IV. PRINCIPLES FOR ANALYSING STRONG-FIELD IONISATION DATA

A. Estimating particle arrivals from detector signals

Charged-particle detectors work by amplifying the electrical current received at the detector input due to the arrival of the particles. Multichannel plates are the most commonly used detectors in strong-field ionisation experiments, though channel electron multipliers are also used. Since most signal-processing electronics uses voltage rather than current, one normally converts the detector output current to a voltage using a resistor or a transimpedance amplifier.

Present-day detectors amplify so strongly that the detection of a single particle results in a readily detectable voltage pulse. However, a certain fraction of particle arrivals simply fail to register at the detector at all. The fraction of arrivals that result in a voltage pulse is called the "quantum efficiency" of the detector. It is straightforward to ensure that the quantum efficiency is independent of particle momentum. On the other hand, when no particles are

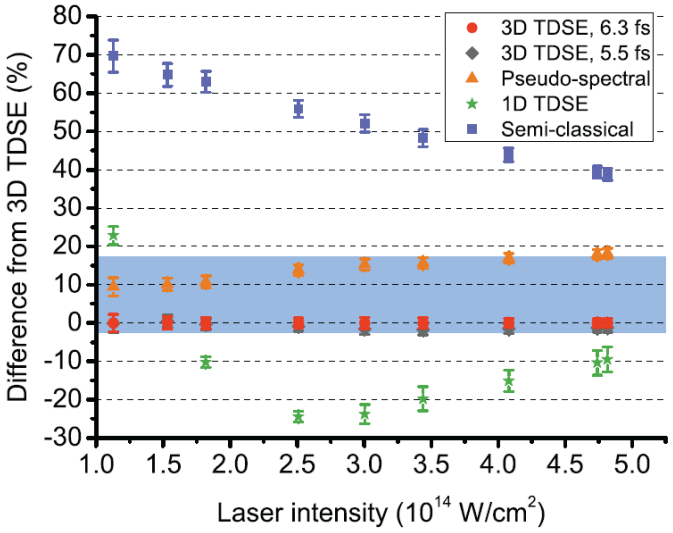

Figure 7. Comparison of theoretical predictions for use in intensity calibration. For each theory, the fractional difference ϵ_X from the 6.3 fs 3D TDSE prediction is shown as a function of the intensity retrieved by the 6.3 fs 3D TDSE. See Eq. (7) for definition of ϵ_X. Gray bar: uncertainty in intensity for independent measurements of laser beam parameters.

arriving, the voltage fluctuations are much smaller than the typical pulse resulting from a particle detection.

Detectors are commonly operated either in counting mode or in proportional mode. In counting mode, the successful detection of a particle gives rise to short voltage pulses on all the readouts V_j. As long as the time between detection events is much longer than the voltage pulse duration $\tau_{\rm det}$, each detection event can be readily distinguished and "counted". In proportional mode, one integrates the entire voltage waveform $V_j(t)$ over an appropriate time interval. The integral is nominally proportional to the number of detected particles.

Counting-mode detection is highly immune to electronic noise. For instance, suppose we wish to count the number of electrons released by a laser pulse in order to make a photoelectron yield measurement. We set up a simple detector with a single voltage readout $V(t)$, which rises to a peak voltage V_0 when a particle is successfully detected. To count particles, we count the number of times that $V(t)$ rises above a fixed discriminator voltage $V_{\rm disc} < V_0$. Whenever this happens, we register a "click", or counting event. Voltage noise of amplitude σ_V is rejected when $\sigma_V \ll V_{\rm disc} \ll V_0 - \sigma_V$. In practice, voltage noise has a negligible effect on counting-mode data whenever reasonable care is taken with the electronics: all signals carried by short coaxial cables, all circuits mounted in sealed metal cases, and all electronics powered by low-noise power supplies.

In counting mode, at most one click can be registered per time interval $\tau_{\rm det}$. If two particles are detected within the same time interval, the detector simply generates a larger voltage pulse within the same time interval. Only one click is generated, even though two particles have arrived. One could imagine setting another, higher discriminator voltage to account for two-particle arrivals. Unfortunately, this tactic fails because the amplitude of the detector pulses varies randomly from pulse to pulse over a wide range - as large as a factor of two. A second discriminator would frequently mis-identify single arrivals as doubles and vice versa.

Whenever the counting-mode detection rate becomes large compared to $1/\tau_{\rm det}$, we must account for the possibility of two-particle detection events being counted as a single click. The particle arrivals are uncorrelated within the time interval $\tau_{\rm det}$, so the event distribution is Poissonian. Writing α for the average number of detection events within $\tau_{\rm det}$, we see that the probability of k events is given by

$$P_k(\alpha) = \frac{e^{-\alpha}\alpha^k}{k!} \tag{8}$$

We *fail* to observe a click in the interval $\tau_{\rm det}$ only if the number of successful detection events is zero, so the probability of observing a click is

$$P_{\rm click} = 1 - P_0(\alpha) = 1 - e^{-\alpha} \tag{9}$$

We notice that $P_{\rm click} \approx \alpha$ for $\alpha \ll 1$ and that $P_{\rm click} \to 1$ as $\alpha \to \infty$. The average number of events is then

$$\alpha = -\ln(1 - P_{\rm click}) \tag{10}$$

Even for $P_{\rm click} = 0.5$, we already have $\alpha = 0.69$, so that the approximation $P_{\rm click} \approx \alpha$ is wrong by 40%.

Counting detection becomes unreliable as $P_{\rm click}$ approaches 1, since the error σ_α in our estimate of α becomes much larger than the error in the observed click probability $\sigma_{P_{\rm click}}$. Basic error propagation shows that [27]

$$\sigma_\alpha = \frac{1}{1 - P_{\rm click}} \sigma_{P_{\rm click}} \tag{11}$$

so that, for instance, the estimated error in α is ten times the error in $P_{\rm click}$ for $P_{\rm click} = 0.9$. Since it is difficult to obtain a reliable estimate for $\sigma_{P_{\rm click}}$ without a long data series, we should be cautious about applying Eqs. (10) and (11) for values of $P_{\rm click}$ close to 1.

Proportional-mode detection avoids the saturation effect of counting mode, so it becomes useful when high detection rates are required. If one wishes to measure the rate of some process over many orders of magnitude, as in experiments measuring total photoelectron or photoion yield, proportional mode is almost mandatory: in counting mode, the total measurement time becomes impractically long at the lowest rates. However, in proportional mode, it becomes difficult to distinguish between signal fluctuations and noise, and the details of the voltage waveform become important. Since counting mode is preferred for accurate measurements, we do not discuss the issues of proportional mode further in this review.

B. Data accumulation and long-term noise

For high-accuracy measurements, one almost always accumulates data for a time that far exceeds the interval between laser pulses. Accumulating data over many laser pulses, or "shots", improves the signal-to-noise ratio of the measured particle distribution and thus our ability to estimate the "true" distribution. Estimating the measurement uncertainty therefore requires statistical analysis of the data accumulation process.

Real experimental data is discrete: for a given particle arrival, the particle momentum is determined to lie within a finite volume element of momentum space. In other words, the momentum component along x is measured as lying in the interval $(p_{\rm det,x}, p_{\rm det,x} + \Delta_{\rm det,x})$, and similarly for the momenta along y and z. For simplicity, we denote this interval as $(\vec{p}_{\rm det}, \vec{p}_{\rm det} + \vec{\Delta}_{\rm det})$. The detector resolution often varies with momentum, so in general $\vec{\Delta}_{\rm det} = \vec{\Delta}_{\rm det}(\vec{p}_{\rm det})$ is a function of momentum. Upon the jth laser shot, we obtain the data $d_j(\vec{p}_{\rm det})$, where $d_j(\vec{p}_{\rm det})$ is the (integer) number of particles that is detected in the momentum interval $(\vec{p}_{\rm det}, \vec{p}_{\rm det} + \vec{\Delta}_{\rm det})$.

By accumulating data over $j = 1, \ldots N$ laser shots, we obtain $D_j(\vec{p}_{\rm det}) = \sum_{j=1}^{N} d_j(\vec{p}_{\rm det})$, where $D_j(\vec{p}_{\rm det})$ is the total (integer) number of particles in the momentum interval. We can therefore estimate the average detection rate within a given momentum interval as $\bar{d}(\vec{p}_{\rm det}) = D_j(\vec{p}_{\rm det})/N$. If every laser shot is identical, the random error in determining $\bar{d}$ can be estimated as [27]

$$\sigma_d(\tau) \approx \frac{\sigma_D}{\sqrt{N}} = \frac{\sigma_D}{\sqrt{\tau \nu_{\rm rep}}} \tag{12}$$

where σ_D is the random error in determining D_j on each laser shot, τ is the total data accumulation time, or "integration time" and $\nu_{\rm rep}$ is the repetition frequency of the laser pulses. Eq. (12) is closely related to the well-known "shot noise limit", and becomes identical to it if $\sigma_D \approx \bar{d}^{1/2}$, i.e., when the D_j follow a Poissonian distribution.

Regrettably, the conditions of the laser shots rapidly cease to be identical in real experiments. It is impractical to continuously monitor every relevant parameter and model its effects on the measurement result. Long-term drift of the experimental parameters begins to affect precision measurement results on the few-second timescale. Hence, the random error of the average value $\bar{d}(\vec{p}_{\rm det})$ rises above the shot-noise value given by Eq. (12). The uncertainties of measurements with integration times $\gtrsim 50$ s are far above the shot-noise value. These considerations also apply to comparisons between measurements taken at different times, e.g., when comparing electron spectra obtained at

different laser intensities. Since the laser parameters continuously drift during the entire series of measurements, the random error in, e.g., the difference between two measurements can easily be much larger than the random error in each measurement.

C. Uncertainty in a single measurement - Allan deviation

We seek to quantify how the random error in the accumulated data, $\sigma_d(\tau)$, depends on the total integration time τ. The Allan deviation is a robust statistical tool for this purpose that was originally developed for characterising clocks and frequency standards [28, 29] but is now used in many areas of metrology and precision measurement [30, 31]. The Allan deviation provides an estimate of $\sigma_d(\tau)$ through the formula

$$\sigma_d^2(\tau) = \frac{1}{2(m-1)} \sum_{j=1} m \left(\bar{d}_{j+1}(\tau) - \bar{d}_j(\tau)\right)^2 \tag{13}$$

where $\bar{d}_j$ is the average value measured between the time $j\tau$ and the time $j(\tau+1)$. There are several refinements of the Allan deviation, e.g., the modified Allan deviation and the overlapping Allan deviation: see [10] and references therein for a more complete discussion. Free software for computing the Allan deviation and its variants is widely available, e.g., as a MATLAB routine [32].

The Allan deviation gives us a simple experimental prescription for estimating the random error $\sigma_d(\tau)$ for any proposed integration time τ. We should take many repeated measurements of $\bar{d}$, with each measurement being integrated over the smallest practical time interval $\tau_{\min}$, and the total time series of measurements spanning a time $\tau_{\max}$. Ideally, we should measure D on every laser shot and record the time series for a total time $\tau_{\max}$ that is far longer than any experimentally relevant integration time. Then, applying the Allan deviation formula Eq. (13) to the time-series data, we obtain an estimate of the random error for any integration time $\tau \lesssim \tau_{\max}$.

Figure 8 illustrates the use of the Allan deviation in a strong-field photoionisation experiment. Panel (a) shows a time series of measurements of photoelectron counts taken at the AASF. Each point in the time series represents a measurement accumulated over an interval $\tau_{\min} = 5$ s, while the total time series lasts for $\tau_{\max} = 3600$ s. Panel (b) shows the corresponding error estimate. For short integration times $\tau \lesssim 40$ s, we see that the Allan deviation is proportional to $1/\sqrt{\tau}$, as predicted by the Eq. (12). On this timescale, drift of the experimental parameters has not yet set in. However, for integration times $\tau \sim 60$ s, the Allan deviation no longer decreases owing to drift. In planning an experiment, we would say that there is probably *no point* in integrating each measurement for longer than 40 – 50 s, as we will not improve our measurement precision by integrating for a longer time. Finally, for integration times $\tau \gtrsim 100$ s, the precision of the measurement actually grows *worse*. By no means should we extend the integration time beyond 100 s. This kind of reasoning has been employed extensively and successfully in planning precision measurement experiments at the AASF [9–11, 33].

The Allan deviation also provides clues as to the experimental sources of long-term measurement error. Figure 8(c) shows a time series of measurements for a key experimental parameter: the power of the AASF laser system, as detected by a standard photodiode. These measurements can be converted to an Allan deviation through Eq. (13) in the same way as the photoelectron counts. The results are shown in Figure 8(d). We see that the Allan deviation of the laser power starts to increase for integration times $\tau \sim 100$, the same timescale as for the Allan deviation of the photoelectron counts in Figure 8(b). While this similarity does not prove that the laser power noise definitely causes the photoelectron count noise, it indicates that there may be an effect worth investigating. On the other hand, Figure 8(e,f) show the time series and Allan deviation for the position of the laser beam along the propagation direction of the atomic beam. The Allan deviation continues to decrease for several hundred seconds, so the instability of the laser beam position does not cause the observed instability in the photoelectron count at the ~ 100 s timescale.

D. Uncertainty of measurement comparisons

For comparisons of measurements taken at different times, we have an additional source of random error due to the drift between measurements. The Allan deviation only quantifies the random error for an individual measurement, in which data is collected during the single-measurement integration time τ. However, if two measurements are separated

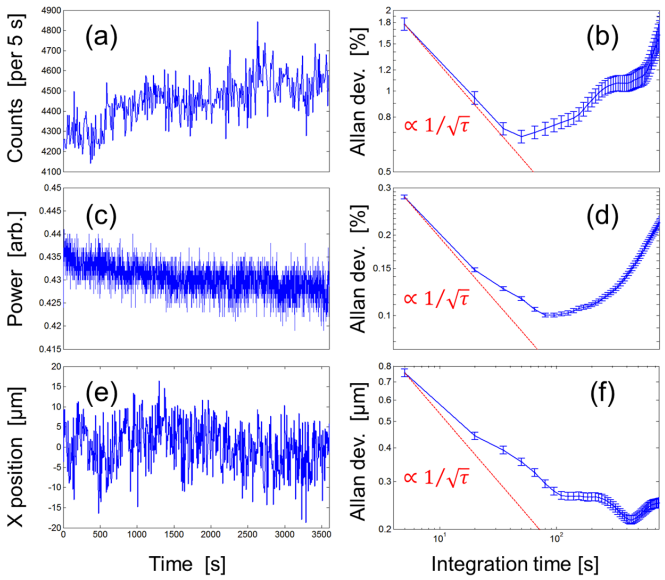

Figure 8. Estimating random measurement error using the Allan deviation. The left column shows the measured time-series data; the right column shows the corresponding Allan deviation. Top row: photoelectron count rate; middle row: laser power; bottom row: laser position along atomic beam axis. Dashed lines in Allan deviation plots show the $1/\sqrt{\tau}$ limit set by Eq. (12). Comparison of the Allan deviations shows that instability of the laser power may limit the useful integration time, while instability of the laser position probably does not.

by a "dead time" $T \gg \tau$, the experimental parameters will still continue to drift during the dead time, and therefore the measurement results will also deviate from each other by an amount that exceeds their individual uncertainties.

One often considers a measurement as a member of a series, for instance, in determining the intensity dependence of photoelectron yield. A conservative error estimate will assume that the dead time between measurements is equal to the entire acquisition time of the series, $T_{\rm ser}$. (Ideally, the measurements are taken in a random order unknown to the experimenters.) The random error of each measurement in the series can then be estimated as

$$\sigma_{d,ser} = \sqrt{\sigma_d^2(\tau) + \sigma_{d,\rm dead}^2(T_{\rm ser})} \tag{14}$$

where $\sigma_{d,\rm dead}(T_{\rm ser})$ is the additional random error arising from the dead-time drift over the time $T_{\rm ser}$.

The dead-time error $\sigma_{d,\rm dead}(T)$ can be quantified from time-series measurements by means of the power spectral density (PSD) of the noise. Just as an optical spectrometer measures the power spectrum of the (time-dependent) electric field, the PSD quantifies the power spectrum of the noise signal [34]. To this end, suppose we obtain a reference time series of m measurements of $\bar{d}$, spanning a total time $\tau_{\rm max}$. The noise PSD can be estimated as [35]

$$S_{dd}(\nu) = \frac{1}{m}\left|\sum_{j=1}^{m} \bar{d}_j e^{-2\pi i j \nu}\right|^2 \tag{15}$$

which is just the squared magnitude of the discrete Fourier transform of the series $\{\bar{d}_j\}$. The variance between measurements of $\bar{d}$ taken at a dead time separation of T is then given by

$$\sigma_{\bar{d},\rm dead}^2(T) = \int_{1/T}^{1/\tau_{\rm min}} d\nu \; S_{dd}(\nu) \tag{16}$$

where $\tau_{\rm min}$ is the time interval for each measurement of $\bar{d}$.

The estimated PSD is often found to be quite noisy for time-series measurements of reasonable duration. However, the PSD of measurements on physical systems almost always follows the form

$$S(\nu) = \begin{cases} A/\nu^{\alpha} & \text{if } \nu \leq \nu_c \quad \text{“flicker noise”} \\ B & \text{if } \nu > \nu_c \quad \text{“white noise”} \end{cases} \tag{17}$$

where A, B, α, and ν_c are constants. By fitting this model to the estimated PSD, one can avoid anomalies in the estimation of $\sigma_{d,\mathrm{dead}}(T)$.

V. PRINCIPLES FOR PREDICTING STRONG-FIELD MEASUREMENT RESULTS

Undertaking an accurate theory-experiment comparison in strong-field ionisation requires careful modelling of the experimental conditions as well as an accurate atomic theory of the strong-field interaction. With generally available experimental techniques, the measured electron distribution cannot be expected to reproduce the output of single-atom simulations.[2] Instead, one measures the distribution that originates from many atoms, each of which is subjected to somewhat different conditions. We therefore need to predict the measured distribution by (1) integrating the single-atom theoretical predictions over variations in the interaction and (2) propagating the integrated distributions through a model of the detector system. The resulting prediction can be expected to match the measurement, but is often quite different from the single-atom theoretical prediction.

A. Theoretical predictions of strong-field ionisation dynamics

First, we remind the reader of what, exactly, constitutes a theoretical single-atom prediction for strong-field ionisation. At nonrelativistic intensities ($\lesssim 10^{16}\,\mathrm{W/cm^2}$) strong-field dynamics in atomic hydrogen is governed by the time-dependent Schrödinger equation (TDSE) for the electron

$$i\hbar\frac{\partial}{\partial t}\psi(\vec{r},t) = \left[\frac{p^2}{2m} - eV(\vec{r}) - e\vec{E}(t)\cdot\vec{r}\right]\psi \tag{18}$$

where $\psi = \psi(\vec{r},t)$ is the electron wavefunction, $V(\vec{r}) = 1/(4\pi\epsilon_0|\vec{r}|)$ is the Coulomb potential, $\vec{E}(t)$ is the laser electric field, and $\vec{p} = -i\hbar\vec{\nabla}$ is the momentum operator.

For atomic H, the exact TDSE can be solved numerically for 800 nm laser wavelength and intensities of 10^{14} to $10^{15}\,\mathrm{W/cm^2}$ by expending substantial time on medium-size computer clusters. The majority of strong-field ionisation experiments fall in this regime. The difficulty of the computation scales rapidly with U_p, so increasing the wavelength or intensity can radically increase computation time. For any atom or molecule with more than one electron or more than one nucleus, the computational problem of the exact TDSE in this regime becomes too difficult, and approximations must be used. As discussed in the Introduction, a major advantage of experiments with atomic H is the opportunity to test these approximations.

The most sophisticated electron detection systems are set up to report the measured momentum $\vec{p}_{\mathrm{det}}$ of the outgoing electrons. Almost all detectors can be considered as measuring some function of the momentum distribution, e.g., the kinetic energy. The laser-atom interaction is long over by the time the electrons are detected, so we can let $t \to \infty$. The electron momentum distribution is proportional to the momentum-resolved ionisation probability,

$$\Pi(\vec{p}) = \lim_{t\to\infty} \left|\langle\vec{p}\,|\,\Psi(\vec{r},t)\rangle\right|^2 \tag{19}$$

$$= \lim_{t\to\infty} \left|\int d\vec{r}\, e^{i\vec{p}\cdot\vec{r}/\hbar}\Psi(\vec{r},t)\right|^2 \tag{20}$$

---

[2] There are prospects for improving this situation. A recent experiment [36] observed the dissociation dynamics of an isolated molecular ion in a UV laser field of $\sim 10^{11}\,\mathrm{W/cm^2}$ intensity.

Various numerical techniques exist for accurately dealing with the $t \to \infty$ limit, so a momentum distribution calculated from an accurate numerical TDSE wavefunction is also accurate.

From the experimental point of view, it is crucial to realise that $\Pi(\vec{p})$ is evaluated for a completely specified laser electric field waveform $\vec{E}(t)$. For the sake of convenience, numerical calculations may use a laser pulse shape that is quite different from the experimentally observed pulse shape. Experimentalists usually characterise the laser field by its intensity, carrier-envelope phase, and pulse shape, with the characterisation of pulse shape presenting special difficulty for few-cycle pulses. All these quantities can fluctuate in time and may also vary spatially depending on the focusing conditions. So far, it appears that the exact pulse shape is relevant for the high-energy plateau electrons of energy $\gtrsim 5U_p$ (see, e.g., [26]), but is not very important at lower energies [10, 16]. However, both the intensity [9, 16] and the carrier-envelope phase [11] are critical parameters over the entire electron energy range.

B. Laser beam geometry

In many experiments, the laser beam can be approximated as a Gaussian beam ([37], ch. 3). For an ideal Gaussian beam propagating along z and focused at $z = 0$, we have

$$I(\rho, z) = I_0 \left(\frac{w_0}{w(z)} \right)^2 e^{-2\rho^2/w^2(z)} \tag{21}$$

$$w(z) \equiv w_0 \sqrt{1 + \left(\frac{z}{z_R} \right)^2} \tag{22}$$

Here I_0 is the peak intensity at the centre of the laser focus, w_0 is the $1/e^2$ intensity radius of the laser spot at the focus, ρ is the distance from the z axis, $z_R = \pi w_0^2/\lambda$ is the Rayleigh range, and λ is the laser wavelength. The parameter $w(z)$ measures the $1/e^2$ spot radius along the propagation axis.

The laser beam geometry also affects the carrier-envelope phase. For a Gaussian beam, the carrier-envelope phase of the pulse evolves according to the Gouy phase

$$\zeta(z) = \tan^{-1} \frac{z}{z_R} \tag{23}$$

The effects of Gouy phase have been directly measured in experiments on strong-field ionisation [38] and high-harmonic generation [33] with noble-gas targets.

The ideal Gaussian beam model may be insufficient when the laser-target interaction region has a significant size along the laser propagation direction. Deviations from an ideal Gaussian beam are especially common for few-cycle pulses, since the spatial mode is easily distorted by hollow-fibre compression and also by misalignment of the reflective focusing optics. While we use the ideal Gaussian beam model in the remainder of this review, Appendix A describes an extended model that is generally adequate to describe non-Gaussian beams encountered in strong-field ionisation experiments. The calculations below can be easily modified to take account of the extended model.

A few experiments in strong-field physics employ special optics to obtain a more uniform intensity distribution than permitted by a Gaussian beam [39, 40]. These techniques show great promise for improving the precision measurement of strong-field ionisation. However, they have not yet been applied to this problem and fall outside the scope of this review.

C. Target interaction geometry

The interaction between the laser and the gas target is determined by the number density of atoms in the target, $n_{\rm at}$. Gas targets naturally fall into three classes: (1) jet targets, in which the gas is emitted from a nozzle, (2) molecular beams, in which a gas jet propagates freely through vacuum and is collimated by apertures, and (3) gas cells, in which the gas uniformly fills a region. Jet targets typically exhibit a smoothly varying, nonuniform density in the interaction region. For molecular beams and gas cells, the density is uniform within the interaction volume,

and zero outside the volume. The laser-target interaction geometries for these target types are depicted in Figure 9. For each type of target, the spot size $w(z)$ of the laser beam typically remains much smaller than the length scale of the molecular density variations.

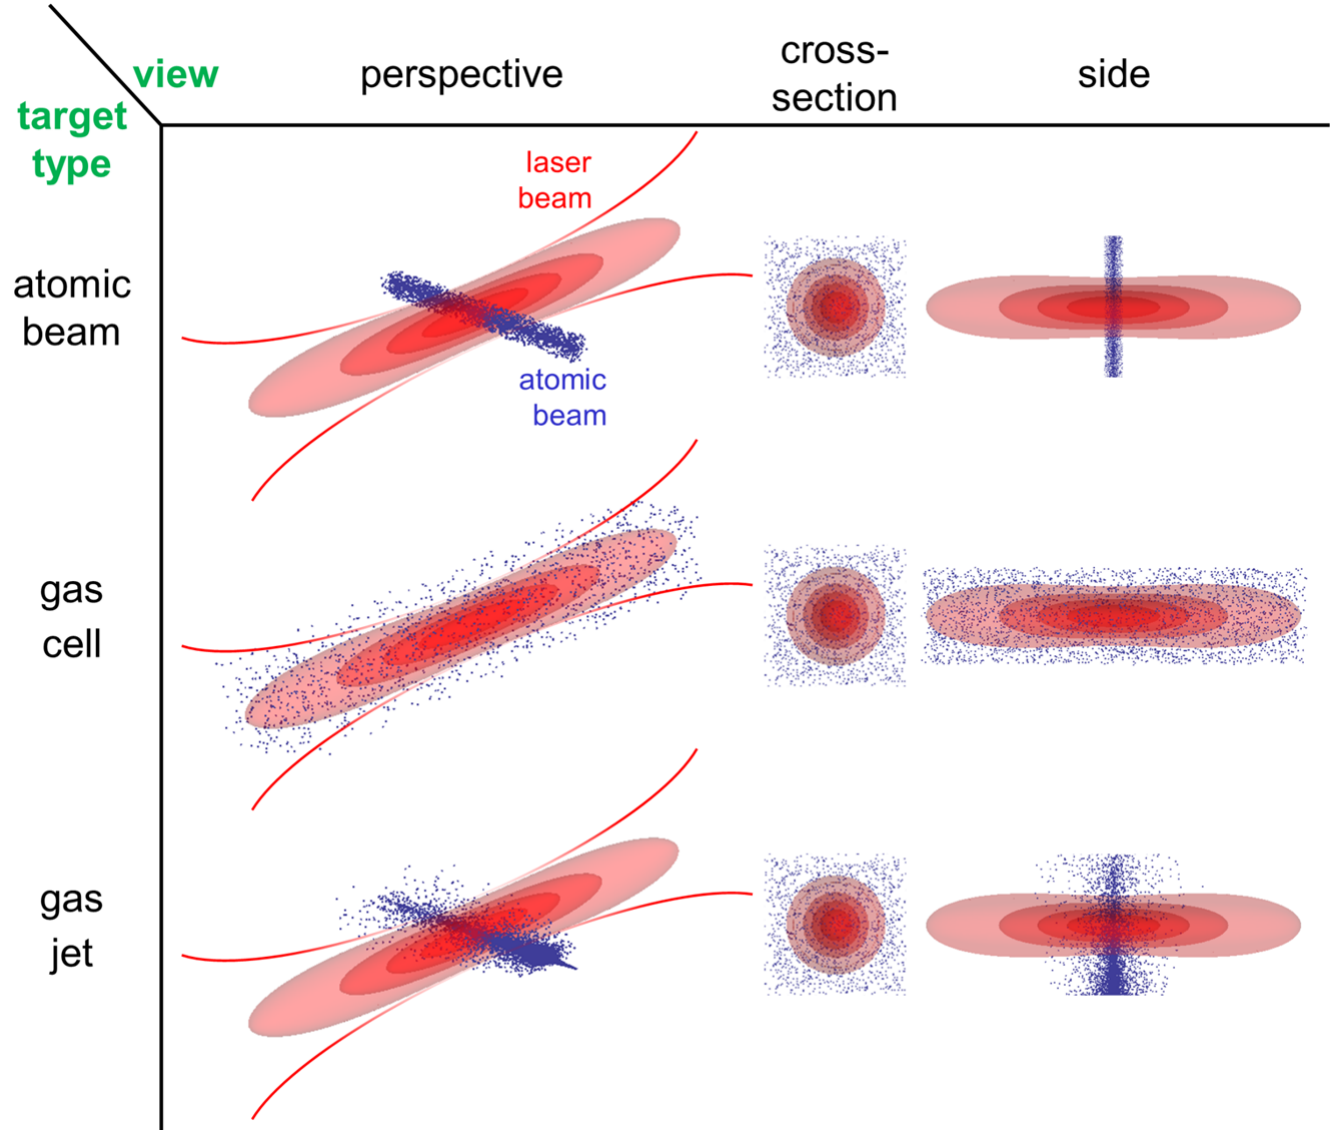

Figure 9. Laser-target interaction geometries for strong-field ionisation. Each row depicts the geometry for one of the three types of gas targets: atomic beam, gas cell, and gas jet. Left column: perspective view of the interaction geometry. Middle column: cross-section in the laser focal plane ($z = 0$). For all three target types, the atomic distribution is independent of the transverse position in the laser beam. Right column: side view perpendicular to both the laser beam and the gas target. Red shading: laser intensity contours; blue dots: atomic distributions; red lines: laser beam edges.

Although gas jets are the most common type of target, the laser-target interaction geometry is poorly defined. In a jet target, atoms are fed into a nozzle from a constant-pressure gas source. The atoms spread out as they emerge from the nozzle, so even if the density is uniform within the nozzle, it quickly assumes a smooth, nonuniform distribution. An effusive jet, with a wide nozzle and low backing pressure, has a $\cos\theta$ angular distribution, where θ is measured from the normal to the nozzle aperture. On the other hand, a narrow nozzle at high pressure will create a supersonic expansion, which can have a very narrow angular distribution.

Atomic beam targets create a sharply defined interaction region by removing the "soft edges" of the gas jet itself. The atomic beam can be made thin relative to the laser Rayleigh range z_R, but wide relative to the laser

spot size $w(z)$. Within the beam, the density is uniform, but outside the beam, it is zero. The atomic beam is formed when a gas jet passes through a high-vacuum flight path containing two thin apertures. The background pressure is kept so low that atoms in the beam do not scatter from background gas in the flight path. Each atom therefore flies along an independent, straight-line trajectory. When the beam encounters an aperture, each atom inside the open part of the aperture passes through undisturbed, but all other atoms are lost from the beam. The diameters of the apertures are typically much smaller than the length of the flight path, so to pass both apertures, a atom must fly straight along the line between them. All atoms emerging from the second aperture are flying parallel to each other, so the density profile of the beam is maintained between the second aperture and the interaction region.

The gas cell is common in experiments with high-energy pulses and/or high-pressure targets. The cell is simply a sealed chamber connected to a constant-pressure gas source. The gas may even just originate in outgassing from the walls of the vacuum chamber. Often the laser beam enters and exits through optical windows. In this case the ionisation rate is guaranteed to be negligible at the entrance and exit; otherwise the windows would immediately be optically damaged. The atomic number density is constant within the cell. Alternatively, the cell may sit in a larger vacuum chamber and the laser beam may enter and exit through small apertures. The number density will then become slightly nonuniform near the apertures.

D. Carrier-envelope phase averaging

Theoretical calculations assume a completely characterised electric field waveform, implying a fixed value of carrier-envelope phase (CEP). Experimentally, the CEP of consecutive pulses from a laser varies randomly over the entire 2π range unless some CEP stabilisation technique is employed. The most common technique measures the CEP with a $f-2f$ interferometer and feeds back the resulting error signal to the laser system, stabilising the CEP [41]. Another active stabilisation technique uses the CEP dependence of strong-field photoelectron spectra to obtain the error signal [42]. One can also use optical parametric amplification to obtain high-energy pulses with passive CEP stability [43]. Nevertheless, it appears challenging to obtain CEP noise below 300 mrad RMS at the laser-atom interaction region, even over timescales of a few seconds [44]. The influence of changing air currents over a several meter beam path, say between the $f-2f$ interferometer and the interaction region, is sufficient to substantially degrade CEP stability.

To compare experiment with theory, we must therefore average the theoretical predictions $\Pi(\vec{p},\phi_{\mathrm{CEP}})$ over the measured distribution of CEP noise. Of course, for a CEP-averaged measurement we must average over a uniform CEP distribution $0 \le \phi_{\mathrm{CEP}} < 2\pi$. We can expand the CEP-dependent momentum distribution in a Fourier series:

$$\Pi(\vec{p},\phi_{\mathrm{CEP}}) = \Pi_0(\vec{p}) + \sum_n \Pi_n(\vec{p}) \sin n(\phi_{\mathrm{CEP}} + \phi_n(\vec{p})) \tag{24}$$

where the coefficients $a_2, a_4, \ldots$ are generally much smaller than $a_1, a_3, \ldots$ [45]. For a Gaussian distribution of CEP noise with standard deviation σ_ϕ, we find that the electron distribution becomes

$$\Pi_{\mathrm{av}}(\vec{p},\phi_{\mathrm{CEP}}) = \int_0^{2\pi} d\phi'\, e^{-\phi'^2/(2\sigma_\phi^2)} \Pi(\vec{p},\phi') \tag{25}$$

$$= \Pi_0(\vec{p}) + \sum_n \kappa_n(\sigma_\phi)\Pi_n(\vec{p}) \sin(\phi_{\mathrm{CEP}} + \phi_n(\vec{p})) \tag{26}$$

$$\text{with } \kappa_n(\sigma_\phi) \equiv e^{-n^2\sigma_\phi^2} \,\mathrm{Re}\left[\frac{\mathrm{erf}\left[(\pi + in\sigma_\phi^2)/(\sqrt{2}\sigma_\phi)\right]}{\mathrm{erf}\left[\pi/(\sqrt{2}\sigma_\phi)\right]}\right] \tag{27}$$

While the phase offsets ϕ_n of the CEP modulation terms are not affected by the noise, the amplitudes are reduced by a factor $\kappa_n(\sigma_\phi)$ that depends sharply on the CEP modulation frequency n. The $\kappa_n(\sigma_\phi)$ are plotted for the first few values of n in Figure 10.

These CEP noise considerations show that precision measurement has a key role to play in revealing the attosecond dynamics of ionisation. For current experiments, with their phase noise of ~ 300 mrad at the interaction region, it appears that only the CEP oscillations with $n = 1$ (period 2π) and $n = 3$ (period $2\pi/3$) are readily accessible. The $n = 3$ oscillation corresponds to a time period of ~ 1 fs at the commonly used 800 nm laser wavelength. Attosecond timing requires us to push to $n \ge 5$, for which the observable CEP effects are an order of magnitude smaller than the

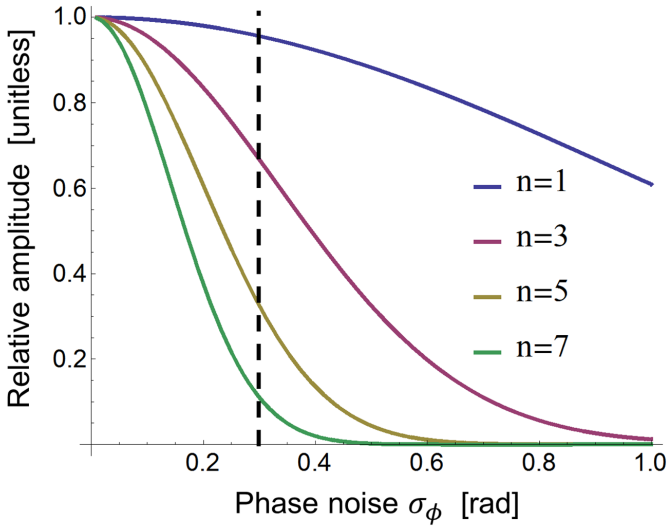

Figure 10. Reduction of CEP effects due to phase noise. The amplitude reduction factors $\kappa_n(\sigma_\phi)$ from Eq. (27) are plotted for the first few values of n. While the lowest-order CEP effect is quite robust to phase noise, the details of CEP variation (higher n) are easily lost for currently accessible levels of phase noise in experiments (vertical dashed line).

generally observed $n = 1$ oscillation effects. Precision measurement techniques become crucial for reliable observations of such small signals.

A particularly common CEP-resolved measurement uses two TOFs on opposite sides of the interaction region to detect high-energy electrons emitted parallel and antiparallel to the laser polarisation direction [46]. This setup is often used as a means of retrieving the CEP [47] and providing a CEP stabilisation signal to the laser system [42]. The two TOFs nominally report the electron spectra $\Pi_U(E_e, \phi_{\rm CEP})$ and $\Pi_D(E_e, \phi_{\rm CEP})$, where E_e is the electron energy. Neglecting detector imperfections, we have $\Pi_D(E_e, \phi_{\rm CEP}) = \Pi_U(E_e, \phi_{\rm CEP} + \pi)$ by symmetry. The CEP effects are often parametrised in terms of the asymmetry:

$$A(E_e, \phi_{\rm CEP}) \equiv \frac{\Pi_U(E_e, \phi_{\rm CEP}) - \Pi_D(E_e, \phi_{\rm CEP})}{\Pi_U(E_e, \phi_{\rm CEP}) + \Pi_D(E_e, \phi_{\rm CEP})} \tag{28}$$

For all but the shortest pulses ($\lesssim 4$ fs), the CEP variations of the sum $\Pi_U + \Pi_D$ are much smaller than the difference $\Pi_U - \Pi_D$. Hence, the amplitude of the asymmetry oscillations at period $2\pi/n$ is reduced by the same factor κ_n given in Eq. (27). For the very shortest pulses, $\Pi_U + \Pi_D$ depends nontrivially on the CEP and there is no simple closed-form expression for the CEP noise effects on the asymmetry.

E. Focal volume integration

The photoelectron yield of strong-field ionisation depends strongly on the intensity, so spatial variations of the intensity within the laser-atom interaction region are a major effect in theory-experiment comparisons. The number density of atoms $n_{\rm at} = n_{\rm at}(\vec{r})$ may also vary within the interaction region. However, as mentioned above, theoretical calculations provide the momentum distribution $\Pi(\vec{p})$ at a constant peak intensity I. The resulting disparity between experiment and theory has been long recognised [48] and was qualitatively discussed for the case of hydrogen in [8]. To include this "focal volume integration" (FVI) effect, we must integrate the momentum distribution over the distribution of laser peak intensities, obtaining

$$S(\vec{p}) = \int d\vec{r}\, n_{\rm at}(\vec{r})\, \Pi(\vec{p}; I(\vec{r})) \tag{29}$$

where $I(\vec{r}) = I(\rho, z)$ is given by the Gaussian beam intensity, Eq. (21). The FVI distribution $S(\vec{p})$ gives the momentum distribution of all electrons (or ions) created in the interaction region.

FVI can also induce additional CEP averaging through the Gouy phase. If the interaction region is significantly extended along the laser propagation direction, the CEP will vary within the interaction region. We have

$$S(\vec{p}, \phi_{\mathrm{CEP},0}) = \int d\vec{r}\, n_{\mathrm{at}}(\vec{r})\, \Pi(\vec{p}; I(\vec{r}), \phi_{\mathrm{CEP}}(z)) \tag{30}$$

where $\phi_{\mathrm{CEP},0}$ is the CEP at $z = 0$ and $\phi_{\mathrm{CEP}}(z) = \phi_{\mathrm{CEP},0} + \tan^{-1}(z/z_R)$ according to Eq. (23).

Atomic beam experiments permit simple and reliable evaluation of FVI effects. The atomic number density is constant inside the beam, $n_{\mathrm{at}} = n_0$, and zero outside. We set the atomic beam diameter d_{beam} so that $w_0 \ll d_{\mathrm{beam}} \ll z_R$. The laser focus intersects the centre of the atomic beam at right angles. The interaction region is approximately a cylinder with diameter w_0 and length d_{beam}. The faces of the cylinder are flat, since $w_0 \ll d_{\mathrm{beam}}$. The laser spot size does not change appreciably in the interaction region since $d_{\mathrm{beam}} \ll z_R$. The FVI equation (29) takes on the "2D" form

$$S_{\mathrm{beam}}(\vec{p}) = 2\pi d_{\mathrm{beam}} n_0 \int_0^\infty \rho\, d\rho\, \Pi\left(\vec{p}; I_0 e^{-2\rho^2/w_0^2}\right) \tag{31}$$

If the laser focus is in the centre of the atomic beam, the Gouy phase shift over the interaction region can be estimated as $\Delta\phi_{\mathrm{Gouy}} \approx \pi d_{\mathrm{beam}}/z_R$. The effect on CEP-resolved measurements is negligible as long as $\Delta\phi_{\mathrm{Gouy}}$ is small compared to the CEP noise $\sigma_\phi \gtrsim 300$ mrad. In practice, one therefore requires $d_{\mathrm{beam}} < z_R/10$.

FVI can also be straightforward for gas cell experiments, since the atom number density can be assumed constant everywhere inside the cell. If the laser beam conforms to the ideal Gaussian model of Sec. III A 2, the FVI equation (29) simplifies to the "3D" form

$$S(\vec{p}) = 2\pi n_0 \int_{z_{\mathrm{in}}}^{z_{\mathrm{out}}} dz \int_0^\infty \rho\, d\rho\, \Pi(\vec{p}; I(\rho, z)) \tag{32}$$

with $z_{\mathrm{in,out}}$ being the positions of the cell entrance and exit, measured relative to the position of the laser focus, and $I(\rho, z)$ given by Eq. (21). In this case, one can approximate $z_{\mathrm{in}} \to -\infty$ and $z_{\mathrm{out}} \to \infty$, and the simplified FVI equation (32) should be extremely accurate. As mentioned above, gas cells that use pinholes for entrances and exits will experience density variations near the pinholes, degrading the accuracy of Eq. (32).

Even though the mathematics of FVI is simple for both gas cells and molecular beams, the data from molecular beams is easier to interpret than the data from gas cells. Features like the above-threshold ionisation plateau are already harder to see in molecular beam data than in theoretical predictions because of FVI over the 2D laser slice. Since gas cells incur additional FVI along the laser propagation direction, such features will be washed away further. In CEP-resolved measurements, the Gouy phase shift along the propagation direction also causes substantial CEP averaging:

$$S(\vec{p}, \phi_{\mathrm{CEP},0}) = 2\pi n_0 \int_0^\infty \rho\, d\rho \int_{z_{\mathrm{in}}}^{z_{\mathrm{out}}} dz\, \Pi(\vec{p}; I(\rho, z), \phi_{\mathrm{CEP},0} + \zeta(z)) \tag{33}$$

The more complex FVI for gas cells also requires a much more careful evaluation of the laser beam quality, which can easily depart from the ideal Gaussian model. If significant laser beam aberrations are present, one must use the more complex laser model of Appendix A in conjunction with the full FVI equation (29) to evaluate the FVI.

There is no straightforward way to evaluate FVI for gas jets. As discussed above, a typical gas jet exhibits a smooth variation of $n_{\mathrm{at}}(\vec{r})$ that is not known *a priori*. If the gas jet diameter becomes significant compared to z_R, the details of $n_{\mathrm{at}}(\vec{r})$ become important. Reliable FVI will then require a careful simulation of the gas dynamics or auxiliary experimental characterisation of the 3D number density. It will also be necessary to carefully characterise the laser beam using the model of Appendix A.

To simplify FVI for a gas jet target, one can block the electrons (or ions) that originate outside a restricted interaction volume [49]. A slit is placed near the interaction region, between the interaction region and the detector. Suppose the electrons fly in a straight line to the detector, and only the particles that pass the aperture reach the detector. Working backward along the line of sight from the detector, we see that the slit defines a sharp slice of the interaction region. FVI becomes simplest for a long slit, with length perpendicular to the laser propagation direction and width $d_{\mathrm{slit}} \ll z_R$. In this case, the FVI reduces to the same form as for the molecular beam, Eq. (31), with d_{beam}

replaced by $d_{\rm slit}$.

For ultimate accuracy, the atomic beam is probably preferable to the arrangement of the gas jet plus slit. In the latter case, one may worry about the effect of any stray electric or magnetic fields in the vicinity of the interaction region, which would break the assumption that the ionised particles travel in a straight line through the slit. The apertures in the atomic beam only need to block neutral particles, which are relatively unaffected by stray fields. The gas jet density should also be low enough that the ionised particles do not scatter from the neutral gas as they emerge; a properly operating atomic beam generally satisfies this condition, since the beam is thin and the density is in the molecular flow regime. However, we are not aware of any experiments that have rigorously compared the accuracy of the two arrangements.

F. Detector system modeling

Real electron detector systems never have perfect momentum resolution. In fact, no detector system measures momentum directly - the raw data consists of electronic pulses that register the arrival time and/or arrival position at a multichannel plate or channel electron multiplier. Knowledge of the geometry and electron-optical design of the detector system then permits one to estimate the properties of the momentum distribution.

In the most common electron detection system, the time-of-flight spectrometer (TOF), the photoelectrons leave the interaction region, propagate freely through a field-free space, and arrive at a microchannel plate. Dividing the field-free path length by the measured time of flight gives the velocity v along the detection direction. The solid angle of electron collection is determined by the path length, the microchannel plate diameter, and sometimes, apertures in the flight path.

While TOF timing measurements can be made quite accurate, much more work is required to ensure that the inferred electron velocity is equally accurate. The length of a given electron trajectory is not simply equal to the physical path length and the electron velocity is not necessarily constant. Stray electric and magnetic fields can easily distort the electron trajectory. In many cases, electron optics are used near the interaction region to increase the solid angle of collection, and their fields also change the trajectory. The effects of these fields depend on the initial electron velocity and direction. Hence, the time of flight can deviate from the expected $1/v$ dependence and the solid angle Ω also depends on v.

The TOF is not a special case - we have used it as an example to illustrate similar systematic effects that are present in all electron detection systems. More generally, our detection system gives us some raw data, e.g., some time-dependent voltage readouts $V_j(t)$, which are subject to noise processes. Our task is to estimate the electron momentum distribution $S(\vec{p})$ at the interaction region. The detection system may also have some auxiliary settings W_k that we can change in order to get more information. The field of estimation theory deals with methods for performing this task. Fortunately, all detection systems share some common features that simplify the problem.

To translate the predicted electron distribution from the interaction region to the detector, we define the detector function $\kappa(\vec{p}_{\rm det}|\vec{p},\{V_k\})$. The detector function is equal to the probability that a particle is observed at momentum $\vec{p}_{\rm det}$, conditional on the particle having initial momentum $\vec{p}$, and given the auxiliary detector parameters $\{V_k\} = V_1,\ldots,V_n$. (This notation conforms to the standard probability notation $P(A|B)$, read as "the probability of observing A conditional on the given facts B".) The auxiliary parameters might include, e.g., the voltages applied to electron-optical elements. Once the detector function is known, we can predict the momentum distribution observed by the detector:

$$Y(\vec{p}_{\rm det};\{V_k\}) = \int d\vec{p}\,\kappa(\vec{p}_{\rm det}|\vec{p},\{V_k\})S(\vec{p}) \tag{34}$$

where $S(\vec{p})$ is the focal-volume-integrated momentum distribution. For detectors that do not resolve the full momentum distribution, the measured momentum $\vec{p}_{\rm det}$ should be replaced with the measured property, e.g., the electron energy in the case of a TOF detector.

Detector systems are generally cylindrically symmetric about an axis $\hat{\zeta}$, and the so-called energy-resolving detectors, such as TOFs or repellers, actually measure the momentum component p_ζ along the detector axis. (In the literature, the detector axis is usually labeled $\hat{z}$. Here we denote it by $\hat{\zeta}$ to avoid confusion with the laser optic axis $\hat{z}$.) Ideally,

the angle θ between $\vec{p}$ and the $\hat{\zeta}$ axis should satisfy $\theta \ll 1$ for all the detected electrons. The energy E can then be written as

$$E \approx \frac{p_\zeta^2}{2m}\left(1 + \frac{1}{2}\theta^2 + \dots\right) \tag{35}$$

If the range of θ is sufficiently small, we can neglect θ entirely and write $p_\zeta = \sqrt{2mE}$. Writing ϕ for the azimuthal angle about the $\hat{\zeta}$ axis, the momentum distribution $\Pi(\vec{p})$ can be rewritten as

$$\Pi(E, \theta, \phi) = \sqrt{\frac{m}{p_\zeta}}\Pi(\vec{p}) \tag{36}$$

where the proportionality constant arises from the change of variables. For a general setting of the detector axis $\hat{\zeta}$, the momentum distribution $\Pi(E, \theta, \phi)$ will depend on all three variables E, θ, ϕ, even though the momentum distribution may often be cylindrically symmetric in some other coordinates.

It is desirable to design the detector system so that the detector function can be approximated as the product of an energy acceptance function and an angular acceptance function

$$\kappa(E_{\rm det}|\vec{p}, \{V_k\}) \approx \kappa_E(E_{\rm det}; E, \{V_k\})\Theta(\theta_{\rm max} - \theta) \tag{37}$$

where E, $E_{\rm det}$ are the electron energy at the interaction region and the electron energy read out by the detector, $\theta_{\rm max}$ is the detector acceptance angle, and $\Theta(x)$ is the step function defined by $\Theta(x < 0) \equiv 0$ and $\Theta(x \geq 0) \equiv 1$. The electron energy spectrum measured by the detector becomes

$$Y(E_{\rm det}; \{V_k\}) = \int_0^{2\pi} d\phi \int_0^{\theta_{\rm max}} d\theta \int_0^{\infty} dE\ \sin\theta\, \kappa_E(E_{\rm det}; E, \{V_k\})\Pi(E, \theta, \phi) \tag{38}$$

In the general case, the dependence of Π on ϕ prevents further simplification of the angular integral.

For an ideal TOF, operating at its default values of the auxiliary parameters $\{V_k\}_{\rm def}$, the energy acceptance function can be approximated by the Dirac delta function

$$\kappa_E(E_{\rm det}; E, \{V_k\}_{\rm def}) = \delta(E_{\rm det} - E) \tag{39}$$

Ideally, the acceptance angle of the TOF, $\theta_{\rm max}$, is almost zero. Then $\Pi(E, \theta, \phi)$ becomes almost independent of θ and ϕ. If all these stringent assumptions are satisfied, the detected energy spectrum assumes the simple form

$$Y(E_{\rm det}; \{V_k\}_{\rm def}) = \pi\theta_{\rm max}^2\Pi(E, 0, 0) \tag{40}$$

In reality, the TOF will exhibit some finite energy resolution, which can be determined by electron trajectory simulations.

For an ideal repeller detector, there is a single auxiliary parameter V_D of central importance. For a given setting of the deflection voltage V_D, the detector measures a single number Y. The energy acceptance function can be approximated as

$$\kappa_E(E_{\rm det}; E, V_D) = \Theta(E - E_c(V_D)) \tag{41}$$

where the electron cutoff energy E_c is determined by V_D. If we again assume that the acceptance angle $\theta_{\rm max}$ is almost zero, the ideal repeller detector therefore measures

$$Y(V_D) = \pi\theta_{\rm max}^2 \int_{E_c(V_D)}^{\infty} dE\, \Pi(E, 0, 0) \tag{42}$$

and inverting the relationship $E_c = E_c(V_D)$ gives us the yield as a function of cutoff energy, $Y(E_c)$.

In reality, the approximation $\theta_{\rm max} \to 0$ is too stringent: the detection rate falls to zero in this limit. Also, the energy acceptance function is never perfectly sharp. To make accurate predictions of the data in Section III, it was necessary to estimate the full detection function using simulations of the electron trajectories. These simulations were carried out using the commercial software package `SIMION` . Figure 11 shows the numerically simulated relationship between the deflection voltage, the energy cutoff, and the acceptance angle. The line between blue and red regions marks the approximate relationship $E_c(V_D) = 0.8(V_D - 5)$ discussed in Section III A 3. The detected yield was calculated using this relationship in conjunction with Eq. (34). Approximating the acceptance function with a perfectly sharp cutoff led to significant underestimates of the yield at low cutoff energy.

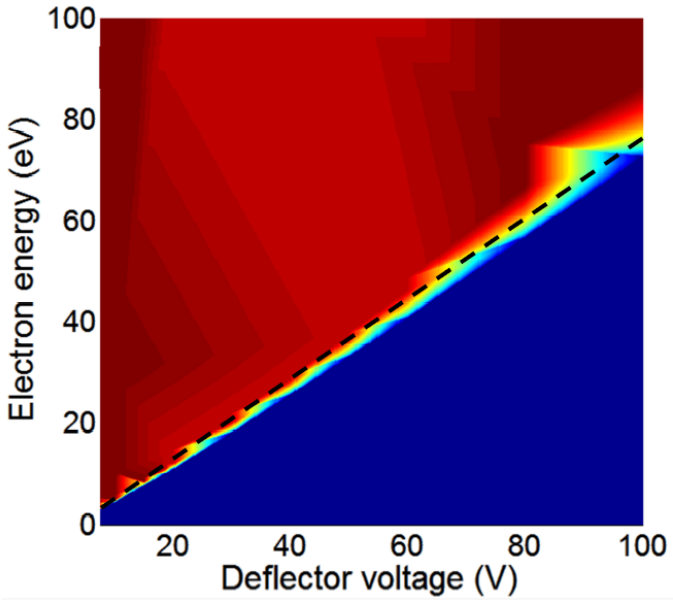

Figure 11. Numerical simulation of the acceptance angle for the repeller detector described in Section III A 3. The colors represent the value of the acceptance angle, from blue (0 mrad) to red (120 mrad). The dotted line marks the approximately linear relationship between cutoff energy and deflection voltage.

VI. SUMMARY AND OUTLOOK

Experiments on atomic hydrogen have provided important benchmarks for strong-field ionisation studies, from qualitative interpretations of multiphoton ionisation and the photoelectron plateau, to quantitative theory-experiment comparisons and calibration of laser parameters. Present-day measurements can resolve the differences between exact 3D TDSE simulations and approximate predictions, even for common and reasonable approximations. The achievement of experimental agreement with the 3D TDSE has validated a concrete path to accurate photoelectron measurements with uncertainties at the 10% level. The methods outlined here are sufficient to find accurate estimates of experimental uncertainty and to convert raw simulation results into accurate predictions of experimental measurements. The laser peak intensity, a key physical parameter for strong-field ionisation, can be accurately retrieved from experimental measurements with 1% accuracy.

Since H is unlikely to be conveniently available in the near future, we envisage that it will serve as a primary reference standard from which secondary calibration standards can be derived. By reaching experiment-theory agreement for H at a given level of uncertainty, one demonstrates the absence of systematic errors at that level. Substituting other atomic and molecular species in the same apparatus, and using otherwise identical experimental techniques, one can obtain new data that is free of systematic errors. By choosing common reference species like the noble gases or simple molecules, one can also obtain certified reference data with a calibration traceable to H. The reproduction of the reference data in another experiment, in another laboratory, then certifies the accuracy of that experiment without the need to access H directly. Such reference data also poses a stringent test of approximate theories in cases where these theories are the only ones available.

We have only begun to explore the potential applications of experiments on H. Measurements like those described here could provide laser intensity calibration at the mid-IR wavelengths characteristic of recent attosecond science experiments. In the mid-IR, the difficulty of standard laser intensity measurements is significantly greater than at 800 nm. Better CEP-resolved measurements could provide a highly accurate calibration of laser CEP by referencing the absolute CEP to 3D-TDSE theory for H. From the point of view of basic physics, the ultimate goal might be a precision measurement of the full electron-ion kinematics of H ionisation using the reaction microscope technique.

Improvements in measurement precision generally lead to improved control of systems, and we expect that strong-field ionisation will be no exception. Indeed, for the highly nonlinear physics of strong-field ionisation, the ability to compare control outcomes with exact theory is even greater than usual. In the future, H may well prove to be the

ideal testing ground for the control of strong-field processes at the atomic level.

ACKNOWLEDGMENTS

The few-cycle measurements described in Section III represent the achievements of a large experimental and theoretical collaboration. We gratefully acknowledge the outstanding contributions of every team member.

The data was collected by the H experimental group at the AASF: Michael Pullen, William Wallace, Omair Ghafur, Dane Laban, Adam Palmer, and Han Xu. The repeller detector system was constructed by Prof Friedrich Hanne of the Westfälische Wilhelms-Universität Münster. We thank Prof Daniel Kleppner of the Massachusetts Institute of Technology for assistance in setting up the hydrogen source.

Theoretical simulations were provided by a large number of groups: Brant Abeln and Prof Klaus Bartschat at Drake University, working with Prof Alexei Grum-Grzhimailo of Lomonosov Moscow State University, Daniel Wells and Prof Harry Quiney of the University of Melbourne, Prof Xiao-Min Tong of the University of Tsukuba, and Igor Ivanov and Prof Anatoli Kheifets of the Australian National University.

The present review was supported by the Australian Research Council (ARC) Centre of Excellence for Coherent X-Ray Science (CE0561787) and by the US Air Force Office of Scientific Research under FA2386-12-1-4025. D.K. is supported by an ARC Future Fellowship (FT110100513).

Appendix A: Imperfections of the laser beam

In reality, the laser beam is never a perfect Gaussian. For most purposes, it appears sufficient to model the beam as a "simple astigmatic beam", according to the classification of the International Organization for Standardization (ISO) [50]. In this model, the laser spot is taken to be an elliptical Gaussian with its principal axes along x and y. Owing to astigmatism, the foci in the x and y directions may be located at different positions $z_{0,x}$, $z_{0,y}$. Moreover, the divergence of the beam along either x or y may exceed the divergence of an ideal Gaussian beam of the same spot size by the "beam propagation ratios" M_x^2, M_y^2. The intensity becomes ([51], ch. 6; [50, 52])

$$I(\vec{r}) = I_0 \frac{w_{0,x}}{w_x(z)} \frac{w_{0,y}}{w_y(z)} \exp\left[-\frac{x^2}{w_x^2(z)} + \frac{y^2}{w_y^2(z)}\right] \tag{A1}$$

$$w_x(z) \equiv w_{0,x}\sqrt{1 + \left(\frac{M_x^2(z - z_x)}{z_{R,x}}\right)^2} \tag{A2}$$

Here $w_{0,x}$ is the minimum spot half-width along the x axis at the x focus, with the Rayleigh range $z_{R,x} = \pi w_{0,x}^2/\lambda$, and similarly for $w_{0,y}$. When the beam is well approximated as an elliptical Gaussian ($M_x^2 \approx M_y^2 \approx 1$), the Gouy phase is given by

$$\zeta(z) = \frac{1}{2}\left(\tan^{-1}\frac{z - z_x}{z_{R,x}} + \tan^{-1}\frac{z - z_y}{z_{R,y}}\right) \tag{A3}$$

However, we are not aware of a simple formula for the Gouy phase for the general case of the simple astigmatic beam. Iterative phase retrieval can be used to compute the Gouy phase from a suitable set of beam intensity distributions acquired at different points along the propagation axis [33].

Evaluating the beam parameters for the simple astigmatic model is nontrivial. Specifications for commercially available laser systems include information on the beam parameters at the laser output. However, off-axis focusing of the beam by a spherical mirror will introduce further astigmatism. Off-axis paraboloidal mirrors can, in principle, focus without astigmatism, but are highly sensitive to imperfect alignment.

[1] B Sheehy and LF DiMauro, "Atomic and molecular dynamics in intense optical fields," Annu. Rev. Phys. Chem. **47**, 463–494 (1996).

[2] P. B. Corkum and F. Krausz, "Attosecond science," Nature Phys. **3**, 381–387 (2007).

[3] Joint Committee for Guides in Metrology, *JCGM 200:2012. International vocabulary of metrology - Basic and general concepts and associated terms (VIM)*, Geneva, Switzerland (2012).

[4] Joint Committee for Guides in Metrology, *JCGM 100:2008. Evaluation of measurement data – Guide to the expression of uncertainty in measurement*, Geneva, Switzerland (2008).

[5] H Rottke, B Wolff, M Brickwedde, D Feldmann, and KH Welge, "Multiphoton ionization of atomic hydrogen in intense subpicosecond laser pulses," Phys. Rev. Lett. **64**, 404 (1990).

[6] H Rottke, D Feldmann, B Wolff-Rottke, and K H Welge, "Atomic hydrogen in intense short laser pulses: a new series of photoelectron peaks from above-threshold ionization," J. Phys. B **26**, L15 (1993).

[7] H Rottke, B Wolff-Rottke, D Feldmann, KH Welge, M Dörr, RM Potvliege, and R Shakeshaft, "Atomic hydrogen in a strong optical radiation field," Phys. Rev. A **49**, 4837 (1994).

[8] G G Paulus, W Nicklich, F Zacher, P Lambropoulos, and H Walther, "High-order above-threshold ionization of atomic hydrogen using intense, ultrashort laser pulses," J. Phys. B **29**, L249 (1996).

[9] M. G. Pullen, W. C. Wallace, D. E. Laban, A. J. Palmer, G. F. Hanne, A. N. Grum-Grzhimailo, B. Abeln, K. Bartschat, D. Weflen, I. Ivanov, A. Kheifets, H. M. Quiney, I. V. Litvinyuk, R. T. Sang, and D. Kielpinski, "Experimental ionization of atomic hydrogen with few-cycle pulses," Opt. Lett. **36**, 3660–3662 (2011).

[10] M. G. Pullen, *Above threshold ionisation of atomic hydrogen using few-cycle pulses*, Ph.D. thesis, Griffith University, Brisbane, Australia (2011), available at http://libraryguides.griffith.edu.au/theses.

[11] W C Wallace, M G Pullen, D E Laban, O Ghafur, H Xu, A J Palmer, G F Hanne, K Bartschat, A N Grum-Grzhimailo, H M Quiney, I V Litvinyuk, R T Sang, and D Kielpinski, "Carrier-envelope phase effects in above-threshold ionization of atomic hydrogen," New J. Phys. **15**, 033002 (2013).

[12] J. Slevin and W. Stirling, "Radio frequency atomic hydrogen beam source," Rev. Sci. Instrum. **52**, 1780–1782 (1981).

[13] J. P. Schwonek, *A study of a cold atomic hydrogen beam source*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA (1990), available at http://hdl.handle.net/1721.1/37496.

[14] B. P. Lavrov, A. V. Pipa, and J. Röpcke, "On determination of the degree of dissociation of hydrogen in non-equilibrium plasmas by means of emission spectroscopy: I. the collision-radiative model and numerical experiments," Plasma Sources Sci. Technol. **15**, 135 (2006).

[15] B P Lavrov, N Lang, A V Pipa, and J Röpcke, "On determination of the degree of dissociation of hydrogen in non-equilibrium plasmas by means of emission spectroscopy: II. Experimental verification," Plasma Sources Sci. Technol. **15**, 147 (2006).

[16] MG Pullen, WC Wallace, DE Laban, AJ Palmer, GF Hanne, AN Grum-Grzhimailo, K Bartschat, I Ivanov, A Kheifets, D Wells, HM Quiney, XM Tong, IV Litvinyuk, RT Sang, and D Kielpinski, "Measurement of laser intensities approaching 10^{15} W/cm^2 with an accuracy of 1%," Phys. Rev. A **87**, 053411 (2013).

[17] T Rathje, NG Johnson, M Möller, F Süßmann, D Adolph, M Kübel, R Kienberger, MF Kling, GG Paulus, and AM Sayler, "Review of attosecond resolved measurement and control via carrier–envelope phase tagging with above-threshold ionization," J. Phys. B **45**, 074003 (2012).

[18] GG Paulus, F Lindner, DB Milošević, and W Becker, "Phase-controlled single-cycle strong-field photoionization," Phys. Scr. **2004**, 120 (2004).

[19] P Hansch, MA Walker, and LD Van Woerkom, "Resonant hot-electron production in above-threshold ionization," Phys. Rev. A **55**, R2535 (1997).

[20] GG Paulus, F Grasbon, H Walther, R Kopold, and W Becker, "Channel-closing-induced resonances in the above-threshold ionization plateau," Phys. Rev. A **64**, 021401 (2001).

[21] SFJ Larochelle, A Talebpour, and SL Chin, "Coulomb effect in multiphoton ionization of rare-gas atoms," J. Phys. B **31**, 1215 (1998).

[22] I V Litvinyuk, K F Lee, P W Dooley, D M Rayner, D M Villeneuve, and P B Corkum, "Alignment-dependent strong field ionization of molecules," Phys. Rev. Lett. **90**, 233003 (2003).

[23] A. S. Alnaser, X. M. Tong, T. Osipov, S. Voss, C. M. Maharjan, B. Shan, Z. Chang, and C. L. Cocke, "Laser-peak-intensity calibration using recoil-ion momentum imaging," Phys. Rev. A **70**, 023413 (2004).

[24] C Smeenk, JZ Salvail, L Arissian, PB Corkum, CT Hebeisen, and A Staudte, "Precise in-situ measurement of laser pulse intensity using strong field ionization," Opt. Express **19**, 9336–9344 (2011).

[25] S Micheau, Z Chen, AT Le, J Rauschenberger, MF Kling, and CD Lin, "Accurate retrieval of target structures and laser parameters of few-cycle pulses from photoelectron momentum spectra," Phys. Rev. Lett. **102**, 073001 (2009).

[26] Z Chen, T Wittmann, B Horvath, and CD Lin, "Complete real-time temporal waveform characterization of single-shot few-cycle laser pulses," Phys. Rev. A **80**, 061402 (2009).

[27] J.R. Taylor, *An Introduction to Error Analysis*, 2nd ed. (University Science Books, Sausalito, CA, USA, 1996).

[28] D W Allan, "Statistics of atomic frequency standards," Proc. IEEE **54**, 221–230 (1966).

[29] J A Barnes, A R Chi, L S Cutler, D J Healey, D B Leeson, T E McGunigal, J A Mullen, W L Smith, R L Sydnor, R F C Vessot, and G M R Winkler, "Characterization of frequency stability," IEEE Trans. Instrum. Meas. **1001**, 105–120 (1971).

[30] D W Allan, "Should the classical variance be used as a basic measure in standards metrology?" IEEE Trans. Instrum. Meas. **1001**, 646–654 (1987).

[31] T. J. Witt, "Practical methods for treating serial correlations in experimental observations," Eur. Phys. J.: Spec. Top. **172**, 137–152 (2009).

[32] M. A. Hopcroft, available at http://www.mathworks.com.au/matlabcentral/fileexchange/13246-allan (2010).

[33] DE Laban, AJ Palmer, WC Wallace, NS Gaffney, RPMJW Notermans, TTJ Clevis, MG Pullen, D Jiang, HM Quiney, IV Litvinyuk, D Kielpinski, and RT Sang, "Extreme ultraviolet interferometer using high-order harmonic generation from successive sources," Phys. Rev. Lett. **109**, 263902 (2012).

[34] R. Paschotta, *Encyclopedia of Laser Physics and Technology* (Wiley-VCH, Berlin, 2008).

[35] K S Shanmugan and A M Breipohl, *Random Signals: Detection, Estimation, and Data Analysis* (Wiley, 1988).

[36] S Kahra, G Leschhorn, M Kowalewski, A Schiffrin, E Bothschafter, W Fuß, R de Vivie-Riedle, R Ernstorfer, F Krausz, R Kienberger, and T Schaetz, "A molecular conveyor belt by controlled delivery of single molecules into ultrashort laser pulses," Nature Phys. **8**, 238–242 (2012).

[37] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, 2nd ed. (Wiley, Hoboken, NJ, USA, 2007).

[38] F Lindner, GG Paulus, H Walther, A Baltuška, E Goulielmakis, M Lezius, and F Krausz, "Gouy phase shift for few-cycle laser pulses," Phys. Rev. Lett. **92**, 113001–113001 (2004).

[39] W Boutu, T Auguste, O Boyko, I Sola, P Balcou, L Binazon, O Gobert, H Merdji, C Valentin, E Constant, E Mével, and B Carré, "High-order-harmonic generation in gas with a flat-top laser beam," Phys. Rev. A **84**, 063406 (2011).

[40] Y Fu, H Xiong, H Xu, J Yao, B Zeng, W Chu, Y Cheng, Z Xu, W Liu, and S L Chin, "Generation of extended filaments of femtosecond pulses in air by use of a single-step phase plate," Opt. Lett. **34**, 3752–3754 (2009).

[41] A Baltuška, M Uiberacker, E Goulielmakis, R Kienberger, V S Yakovlev, T Udem, T W Hänsch, and F Krausz, "Phase-controlled amplification of few-cycle laser pulses," IEEE J. Sel. Top. Quant. Elec. **9**, 972–989 (2003).

[42] D. Adolph, A. M. Sayler, T. Rathje, K. Rühle, and G. G. Paulus, "Improved carrier-envelope phase locking of intense few-cycle laser pulses using above-threshold ionization," Opt. Lett. **36**, 3639–3641 (2011).

[43] C Vozzi, G Cirmi, C Manzoni, E Benedetti, F Calegari, G Sansone, S Stagira, O Svelto, S De Silvestri, M Nisoli, and G Cerullo, "High-energy, few-optical-cycle pulses at 1.5 μm with passive carrier-envelope phase stabilization," Opt. Express **14**, 10109–10116 (2006).

[44] T. Wittmann, B. Horvath, W. Helml, M. G. Schätzel, X. Gu, A. L. Cavalieri, G. G. Paulus, and R. Kienberger, "Single-shot carrier-envelope phase measurement of few-cycle laser pulses," Nature Phys. **5**, 357–362 (2009).

[45] V. Roudnev and B. D. Esry, "General theory of carrier-envelope phase effects," Phys. Rev. Lett. **99**, 220406 (2007).

[46] GG Paulus, F. Grasbon, H. Walther, P. Villoresi, M. Nisoli, S. Stagira, E. Priori, and S. De Silvestri, "Absolute-phase phenomena in photoionization with few-cycle laser pulses," Nature **414**, 182–184 (2001).

[47] G. G. Paulus, F. Lindner, H. Walther, A. Baltuška, E. Goulielmakis, M. Lezius, and F. Krausz, "Measurement of the phase of few-cycle laser pulses," Phys. Rev. Lett. **91**, 253004 (2003).

[48] S Augst, D D Meyerhofer, D Strickland, and SL Chin, "Laser ionization of noble gases by coulomb-barrier suppression," JOSA B **8**, 858–867 (1991).

[49] P Hansch and L D Van Woerkom, "High-precision intensity-selective observation of multiphoton ionization: a new method of photoelectron spectroscopy," Opt. Lett. **21**, 1286–1288 (1996).

[50] International Organization for Standardization, *ISO 11146-1:2005(E). Lasers and laser-related equipment - Test methods for laser beam widths, divergence angles and beam propagation ratios - Part 1*, Geneva, Switzerland (2005).

[51] A. Yariv, *Quantum Electronics*, 3rd ed. (Wiley, 1987).

[52] A. E. Siegman, "How to (maybe) measure beam quality," in *DPSS Lasers: Applications and Issues*, edited by M. W. Dowley (Optical Society of America, Washington, D.C., USA, 1998) pp. 184–199.